

Ontology is the philosophical discipline which aims to understand how things in the world are divided into categories and how these categories are related together. This is exactly what information scientists aim for in creating structured, automated representations, called 'ontologies,' for managing information in fields such as science, government, industry, and healthcare. Currently, these systems are designed in a variety of different ways, so they cannot share data with one another. They are often idiosyncratically structured, accessible only to those who created them, and unable to serve as inputs for automated reasoning. This volume shows, in a non-technical way and using examples from medicine and biology, how the rigorous application of theories and insights from philosophical ontology can improve the ontologies upon which information management depends.



Distributed in North and South America
by Transaction Books
ISBN 978-3-938793-98-5



Katherine Munn, Barry Smith (Eds.) · *Applied Ontology*

METAPHYSICAL RESEARCH

Edited by

Maria E. Reicher · Johanna Seibt
Barry Smith · Daniel von Wachter

Katherine Munn
Barry Smith (Eds.)

Applied Ontology

An Introduction



Katherine Munn, Barry Smith
Applied Ontology
An Introduction

METAPHYSICAL RESEARCH

Herausgegeben von / Edited by

Uwe Meixner • Johanna Seibt
Barry Smith • Daniel von Wachter

Band 8 / Volume 9

Katherine Munn, Barry Smith

Applied Ontology

An Introduction



ontos

verlag

Frankfurt | Paris | Lancaster | New Brunswick

Bibliographic information published by the Deutsche Nationalbibliothek

The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at <http://dnb.d-nb.de>.



North and South America by
Transaction Books
Rutgers University
Piscataway, NJ 08854-8042
trans@transactionpub.com



United Kingdom, Eire, Iceland, Turkey, Malta, Portugal by
Gazelle Books Services Limited
White Cross Mills
Hightown
LANCASTER, LA1 4XS
sales@gazellebooks.co.uk



Livraison pour la France et la Belgique:
Librairie Philosophique J. Vrin
6, place de la Sorbonne; F-75005 PARIS
Tel. +33 (0)1 43 54 03 47; Fax +33 (0)1 43 54 48 18
www.vrin.fr

©2008 ontos verlag
P.O. Box 15 41, D-63133 Heusenstamm
www.ontosverlag.com

ISBN 978-3-938793-98-5

2008

No part of this book may be reproduced, stored in retrieval systems or transmitted in any form or by any means, electronic, mechanical, photocopying, microfilming, recording or otherwise without written permission from the Publisher, with the exception of any material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use of the purchaser of the work

Printed on acid-free paper
FSC-certified (Forest Stewardship Council)

Printed in Germany
by buch bücher **dd ag**

Table of Contents

Introduction: What is Ontology for? Katherine Munn	7
Acknowledgments	20
1. Bioinformatics and Philosophy Barry Smith and Bert Klagges	21
2. What Is Formal Ontology? Boris Hennig	39
3. A Primer on Knowledge Management and Ontological Engineering Pierre Grenon	57
4. New Desiderata for Biomedical Terminologies Barry Smith	83
5. The Benefits of Realism: <i>A Realist Logic with Applications</i> Barry Smith	109
6. A Theory of Granular Partitions Thomas Bittner and Barry Smith	125
7. Classifications Ludger Jansen	159
8. Categories: <i>The Top-Level Ontology</i> Ludger Jansen	173
9. The Classification of Living Beings Peter Heuer and Boris Hennig	197
10. Ontological Relations Ulf Schwarz and Barry Smith	219
11. Four Kinds of ‘Is_A’ Relation Ingvar Johansson	235

12. Occurrents	
Boris Hennig	255
13. Bioinformatics and Biological Reality	
Ingvar Johansson	285
References	311
Index	329

Introduction: What is Ontology for?

Katherine Munn

If you are reading this, then chances are you are a philosopher, an information scientist, or a natural scientist who uses automated information systems to store or manage data.

What these disciplines have in common is their goal of increasing our knowledge about the world, and improving the quality of the information we already have. Knowledge, when handled properly, is to a great extent cumulative. Once we have it, we can use it to secure a wider and deeper array of further knowledge, and also to correct the errors we make as we go along. In this way, knowledge contributes to its own expansion and refinement. But this is only possible if what we know is recorded in such a way that it can quickly and easily be retrieved, and understood, by those who need it. This book is a collaborative effort by philosophers and information scientists to show how our methods of doing these things can be improved. This introduction aims, in a non-technical fashion, to present the issues arising at the junction of philosophical ontology and information science, in the hope of providing a framework for understanding the essays included in the volume.

Imagine a brilliant scientist who solves a major theoretical problem. In one scenario he scribbles his theory on a beer mat, sharing it only with his drinking companions. In this scenario, very few scientists will have the ability to incorporate this discovery into their research. Even were they to find out that the solution exists, they may not have the resources, time, or patience to track it down. In another scenario our scientist publishes his solution in a widely read journal, but has written it in such a sloppy and meandering way that virtually no one can decipher it without expending prohibitive amounts of effort. In this scenario, more scientists will have access to his discovery, and may even dimly recognize it as the truth, but may only understand it imperfectly. No matter how brilliant our scientist is, or how intricately he himself understands his discovery, if he fails to convey it to the scientific community in such a way that they have ready access to it and can understand it, unfortunately that community will not benefit from what he has discovered. The moral of this story is that the means by which knowledge is conveyed are every bit as important as that knowledge itself.

The authors' goal in producing this book has been to show how philosophy and information science can learn from one another, so as to

create better methodologies for recording and organizing our knowledge about the world. Our interest lies in the representation of this knowledge by *automated* information systems such as computerized terminologies and taxonomies, electronic databases, and other knowledge representation systems. Today's automation of knowledge representation presents challenges of a nature entirely different from any faced by researchers, librarians or archivists of the pre-computer age.

Before discussing the unique challenges posed by automated systems for storing knowledge, we must say a few brief words about the term 'knowledge'. We are not using this term in a sense corresponding to most philosophical theories. What these theories have in common is the requirement that, in order for a belief or a state of mind to count as knowledge, it must connect the person to the truth. That is, a belief or a state of mind counts as knowledge only if its representational content corresponds with the way the world is. Most philosophical theories add the condition that this correspondence must be non-accidental: there must be a causal relation between the belief and its being the case; the person must base the belief on a certain kind of evidence or justification, and so forth (pick your theory).

The sense of 'knowledge' used in information science is more relaxed. Terms such as 'knowledge engineering' and 'knowledge management' do not refer to knowledge in the sense of a body of beliefs that are *apodictically true*, but of a body of beliefs which the scientific community has good reason to believe are true and thus treats in every respect as if they are true. Most researchers recognize that some of these highly justified beliefs are not, in fact, knowledge in the strict sense, since further scientific development could show them to be false. Recognizing this is part of what drives research forward; for part of the goal of research is to cause the number of false beliefs to decrease and the number and nuance of true beliefs to increase. The information stored in automated systems constitutes knowledge in the sense of *beliefs which we have every reason to believe are true, but to which we will not adhere dogmatically should we obtain overruling reasons to believe otherwise*. (We will often use 'information' in the same sense as 'knowledge'.) This approach, called *realist fallibilism*, combines a healthy intellectual humility with the conviction that humans can take measures to procure true beliefs about the world.

So much for 'knowledge'. What does it mean to *store* or *represent* knowledge? (We will use these terms interchangeably.) Say that you have a

bit of knowledge, i.e., a belief that meets all the requirements for knowledge. To store or represent it is to put it into a form in which it can be retained and communicated within a community. Knowledge has been stored in such forms as words, hieroglyphs, mnemonics, graphs, oral tradition, and cave scratching. In all of these forms, knowledge can be communicated, passed on, or otherwise conveyed, from one human being to another.

Automated information systems pose unprecedented challenges to the task of storing knowledge. In the same way that knowledge is represented on the pages of a book by one person and read by another, it is entered into an automated system by one person and retrieved by another. But whereas the book can convey the knowledge to the reader in the same form in which the writer recorded it, automated information systems must store knowledge in forms that can be processed by non-human agents. For computers cannot read or understand words or pictures, so as to answer researchers' queries in the way that the researchers would pose them, or to record their findings as researchers would. Computers must be programmed using explicit codes and formulas; hence, the quality of the information contained in information systems is only as high as the quality of these codes and formulas.

Automated information systems present unique opportunities for representing knowledge, since they have the capacity to handle enormous quantities of it. The right technology enables us to record, obtain, and share information with greater speed and efficiency than ever before, and to synthesize disparate items of information in order to draw new conclusions. There are different sorts of ways in which information systems store knowledge. There are databases designed for storing *particular* knowledge pertaining to, for example, specific experimental results, specific patients treated at a given hospital during a given time period, or specific data corresponding to particular clinical trials. Electronic health record (EHR) systems, used by hospitals to record data about individual patients, are examples of databases which store such particular knowledge. There are also systems designed for storing *general* knowledge. General knowledge includes the sorts of statements found in textbooks, which abstract from particular cases (such as *this* patient's case of pneumonia) and pertain, instead, to the traits which most of those particular cases have in common (such as lung infection, chill, and cough). Systems designed to store general knowledge include controlled vocabularies, taxonomies, terminologies, and so forth. Examples of these

include the Gene Ontology, the Foundational Model of Anatomy, and the Unified Medical Language System Semantic Network.

Ideally, these two types of system will play complementary roles in research. Databases and other systems for storing particular information should be able to provide empirical data for testing general theories, and the general information contained in controlled vocabularies and their ilk should, in turn, provide sources of reference for empirical researchers and clinicians. How better, for example, to form and test a theory about pneumonia than by culling the clinical records of every hospital which has recorded cases of it? How better to prepare for a possible epidemic than by linking the electronic record systems of every hospital in the country to a centralized source, and then programming that source to automatically tag any possibly dangerous trends?

But in order for these goals to be realized, automated information systems must be able to share information. If this is to be possible, every system has to represent this information in the same way. For any automated information system to serve as a repository for the information gathered by researchers, it must be pre-programmed in a way that enables it to accommodate this information. This means that, for each type of input an information system might receive, it must have a category corresponding to that type. Therefore, an automated information system must have a categorial structure readymade for slotting each bit of information programmed into it under the appropriate heading. That structure, ideally, will match the structure of other information systems, to facilitate the sharing of information among them. But if this is to be possible, there must be one categorial structure that is common to all information systems. What should that structure look like?

There are several possible approaches to creating category systems for representing information about the world. One approach, which Smith calls the *term orientation* (see Chapter 4), is based on the observation that researchers often communicate their findings in the form of sentences. What better way to create a category system than to base it on the meanings of the words in those sentences? One problem with this approach is that the meaning of a word often does not remain constant; it may change from context to context, as well as over the course of time. Another problem is that natural language cannot be guaranteed to contain a word which encompasses precisely the meaning one wants to express, especially in scientific disciplines that are constantly making discoveries for which there are not yet established words. Another approach, which is standardly

referred to as the *concept orientation*, attempts to get around these difficulties by substituting words with concepts, seen (roughly) as hypostatizations of the meanings of words into mental entities. In other words, a concept is a word whose meaning has been fixed forever in virtue of being attached to a special kind of abstract thing. Thus, even if some slippage occurs between a word and its original meaning, that meaning will always have a concept to which it adheres. One simple problem with this approach (Smith provides a litany) is that it goes to great lengths to posit a layer of reality – that of concepts – for theoretical purposes only. This raises the question why the structure of the world itself should not be used as a guide to creating categories, an approach known as *realism*. After all, our knowledge is *about* the world, not about concepts.

A major contention against realism is that reality is just too massive, diffuse, or limitless, for human understanding to grasp. There are far more things in the world, and far more kinds of things, than any one person can think or know about, even over the course of a lifetime. Ask one hundred people what the most basic underlying categories of the world are, and you will likely get one hundred different answers. Even scientific disciplines, which reflect not the understanding of one person but of successive groups of people with similar goals and methods, can produce no more than a perspective on one specific portion of reality, to the exclusion of the rest. The object of their study is limited to a specific *domain* of reality, such as the domain of living things for biology or the domain of interstellar-objects for astronomy. Human understanding cannot, either individually or collectively, grasp reality as it is in its entirety; hence, the conceptualist does not expect to be able to represent reality in the categories of automated information systems.

The realist response developed in this volume (particularly Chapters 1, 3, 4, 6, and 7) is this: we can and should understand the existence of multiple perspectives not as a hindrance to our ability to grasp that reality as it is, but as a means by which we can obtain a deeper understanding of it. For, from the fact that there are multiple perspectives on reality alone, it does not follow that none – or only one – of these perspectives is *veridical*, i.e., represents some aspect of reality *as it truly is*.

A perspective is merely the result of someone's coming to cognitive grips with the world. Precisely because reality is so multi-faceted, we are forced to filter out some aspects of it from our attention which are less relevant to our purposes than others. Some of these processes of selection are performed deliberately and methodically. For example, biologists set

into relief the domain of living things, in order to focus their study on traits shared by them which non-living things do not have. Forest rangers set into relief the domain of a specific geographical area and certain specific features, such as marked trails and streams, which they represent in maps for the purposes of navigation. Often, especially among scientists, the purpose of roping off a particular domain is simply to gain understanding of what the entities within it have in common, and of what makes them different from entities in other domains.

The selection of a particular perspective is an act of cognitively *partitioning* the world: drawing a mental division between those things upon which we are focusing and those which fall outside our domain of interest. (Chapter 6 develops a theory of how we partition the world.) Take as an example Herbert, who is a frog. Let us imagine that Herbert is a domain of study unto himself. We thereby cognitively divide the world into two domains: Herbert, and everything else.

Given a partitioning of the world into domains, it becomes possible to create sub-partitions within those domains. Herbert happens to be a frog, in addition to being composed of molecules. Each of these features yields a unique perspective from which Herbert can be apprehended: the coarse-grained level of Herbert as a whole single unit, and the fine-grained level of his molecules. Most of us think of Herbert as a single unit because it is as such that we apprehend him in his terrarium. Although we may know that he is composed of molecules, his molecules are not relevant to our apprehension of him, and so we filter them out. A molecular biologist, on the other hand, may think more about Herbert's molecules than about Herbert as a whole, even though he is aware that those molecules constitute a whole frog. There is only one Herbert that we and the molecular biologist apprehend, but, depending upon our interests and our focus, we may each apprehend him from different *granular perspectives*.

Recognizing that there are multiple veridical perspectives on reality is *not* equivalent to endorsing relativism, the view that *all* perspectives are veridical. Here are two examples of non-veridical perspectives on Herbert: one which views him as a composite of the four complementary elements earth, air, fire, and water; another which views him as an aggregate of cells joined by an aberrant metaphysical link to the soul of Napoleon. The existence of multiple perspectives does not imply that we are unable to grasp reality as it is, and the fact that it is possible to obtain deeper understanding of reality *through* those perspectives does not imply that all perspectives are veridical representations of reality.

This is not to suggest that it is always easy to distinguish veridical perspectives from non-veridical ones. In fact, it is this difficulty which forces responsible ontologists and knowledge engineers to temper their realism with a dose of fallibilism. One of the main ways to determine the likelihood of a perspective's being veridical is to assess its explanatory power, that is, the breadth and depth of the explanations it can offer of the way the world works. The four-element perspective on Herbert seemed plausible to certain people at a certain point in history, precisely because it offered a means of explaining the causal forces governing the world. It seems less plausible now because better means of explanation have been developed.

Each automated information system strives to represent a veridical perspective on that partition of reality about which it stores knowledge. As we have seen, there are features intrinsic to such systems which render them better or worse for fulfilling this goal. A system which is programmed with a structure that corresponds closely to the structure of the granular partition itself is more likely to be veridical; think of the four-element perspective versus the molecular one. An information system with the categories 'earth', 'air', 'fire', and 'water' is less likely to serve as basis for an accurate categorization of Herbert's various components than is a system with such categories as 'cell', 'molecule', and 'organ'.

The best kinds of categories are *natural* in the sense that they bring genuine similarities and differences existing in reality to the forefront (this view is developed in Chapters 7 and 8). Natural category divisions tell us something about how the underlying reality truly is. Thus, it is more likely that knowledge of such naturally existing categories will put us in a position to construct systematic representations of that domain which have some degree of predictive power. If we can predict the way in which entities in a domain will behave under certain conditions, we are better able to understand that domain, interact with it, and gain more knowledge about it.

Hence the realist, who believes that it is possible for humans to obtain knowledge about the world, seeks to find out, as best he can, what the natural categories of reality are. His goal as a knowledge engineer is to create an information system that is structured in a way that mirrors those categories. Such a system will be prepared to receive information about as wide an array of entities as possible. Then, it should represent information by tagging each piece of information as being about something that has certain traits which make that thing naturally distinct from other entities.

Now, there is at least one natural category into which every entity falls: the category of existing things. It follows that there is at least one perspective from which all of reality is visible, one partition in which every entity naturally belongs: the partition of existing things. This partition is admittedly large-grained in the extreme; it does not provide us with more than a very general insight into the traits of the entities it encompasses. But it does provide us with insight into one crucial trait, existence, which they all have in common. It is this partition which constitutes the traditional domain of *ontology*.

Ontology in the most general sense is the study of the traits which all existing things have insofar as they exist. (This is an admittedly airy definition of an abstract notion; see Chapter 2 for elaboration). It is significant that the philosophical term ‘ontology’ has been adopted by the information-science community to refer to an automated representation (taxonomy, controlled vocabulary) of a given domain (a point developed in Chapter 1). We will sometimes use the term ‘ontology’ in this sense, in addition to using the philosophical sense expounded in Chapter 2.

Since there is one trait, existence, which all entities in reality have in common at the most general level, it is reasonable to suppose that there are other traits which some entities have in common at more specific levels. This supposition conforms to our common-sense assumption that some entities are more alike than others. If this is correct, it would suggest that our ability to understand something about reality in its entirety does not stop at the most general level, but continues downward into more specific levels. The challenge for the realist is to devise a means to discern the categorial subdivisions further down the line; this challenge is taken up in Chapter 9.

Clearly, an upper-level system of categorization encompassing all entities would be an enormous step toward the goal of optimal knowledge representation. If all information systems were equipped with the same upper-level category system (sometimes called a *domain-independent formal ontology*), and if this category system did exhaust the most general categories in reality, then it would be possible to share information among systems with unprecedented speed, efficiency, and consistency. The contributions in this book are aimed at this long-term, but worthwhile, goal. Although the methods developed here are intended to be applicable to any domain, we have chosen to limit our focus primarily to the domains of biology and medicine. The reason is that there are particularly tangible benefits for the knowledge representation systems in these domains.

Accordingly, in ‘Bioinformatics and Philosophy’ (Chapter 1), philosopher Barry Smith and geneticist Bert Klagges make a case for the use of applied ontology in the management of biological knowledge. They argue that biological knowledge-management systems lack robust theories of basic notions such as kind, species, part, whole, function, process, environment, system, and so on. They prescribe the use of the rigorous methods of philosophical ontology for rendering these systems as effective as possible. Such methods, developed precisely for the purpose of obtaining and representing knowledge about the world, have a more than two thousand year-old history in knowledge management.

In ‘What is Formal Ontology?’ (Chapter 2) Boris Hennig brings that most general, abstract domain of existing things down to earth. His goal is to help us understand what the more specific categories dealt with in this book are specifications of. The historical and philosophical background he provides will enable us to view formal ontology afresh in the present context of knowledge management. That context is illuminated in Pierre Grenon’s ‘A Primer on Knowledge Management and Ontological Engineering’ (Chapter 3). Grenon draws upon non-technological examples for two purposes: first, to explain the task of knowledge management to non-information scientists; second, to highlight the reasonableness of the view that knowledge management is about representing reality. He provides insight into the task of the knowledge engineer, who is promoted to the post of *ontological engineer* when he uses rigorous ontological methods to systematize the information with which he deals. Finally, Grenon describes some current (worrying) trends in the knowledge-management field, for which he prescribes a realist ontological approach as an antidote.

Some of these trends are elaborated upon in Barry Smith’s ‘New Desiderata for Biomedical Terminologies’ (Chapter 4). Smith chronicles the development of the concept orientation in knowledge management, offering a host of arguments against it and in favor of the realist orientation. In ‘The Benefits of Realism: *A Realist Logic with Applications*’ (Chapter 5) Smith goes on to demonstrate the problem-solving potential of a realist orientation. He does so by developing a methodology for linking sources of particular knowledge (such as databases) with sources of general knowledge (such as terminologies) in order to render them interoperable. This would dramatically improve the speed and efficiency of the information-gathering process as well as the quality of the information garnered. Implementing his methodology would require a global switch to

the realist orientation in knowledge management systems. Arduous as such a switch would be, his example shows the massive benefits that it would proffer.

If we are to reconstruct existing knowledge management systems to reflect a realist orientation, we will need a theoretical blueprint to guide us. We must start by formalizing the most basic commitment of the realist orientation, *realist persepectivalism*, which is the view that we can obtain knowledge of reality itself by means of a multiplicity of veridical granular partitions. Bittner and Smith (Chapter 6) provide a formal theory of granular partitions for configuring knowledge management systems to accommodate the realist orientation. Only such a theory, they claim, can provide the foundation upon which to build knowledge management systems which have the potential to be interoperable, even though they deal with different domains of reality.

How do we build up an information system that succeeds at classifying the entities in a given domain on the foundation of a theory of granular partitions? In ‘Classifications’ (Chapter 7), Ludger Jansen provides eight criteria for constructing a good classification system, complete with real examples from a widely used information system, the National Cancer Institute Thesaurus (NCIT), which fails to meet them. Nonetheless, he points out, there are numerous practical limitations which an ontological engineer must take into account when constructing a realist ontology of his domain. Since a classification system is, to some extent, a model of reality, the more limited the knowledge engineer’s resources (temporal, monetary, technological, and so forth), the greater his system must abstract from the reality it is supposed to represent. But the existence of such practical limitations does not require us to abandon the goal of representing reality. Jansen recommends meeting practical needs with accuracy to reality by distinguishing between two types of ontologies with distinct purposes. The purpose of *reference ontologies* is to represent the complete state of current research concerning a given domain as accurately as possible. Alternatively, the purpose of *application ontologies*, such as particular computer programs, should be to fit the most relevant aspects of that information in an application designed with certain practical limitations in mind. Reference ontologies should serve as the basis for creating application ontologies. This way, accuracy to reality can stand side by side with utility without either one needing to be sacrificed. Further, application ontologies that are based on the same reference ontologies will be more

easily interoperable with each other than application ontologies based on entirely different frameworks.

In ‘Categories: *The Top-Level Ontology*’ (Chapter 8), Jansen applies the criteria for good classification to the question of what the uppermost categories of a reference ontology should be. Once we move below the most general category, ‘being’, what are the general categories into which all existing things can be exhaustively classified? Jansen answers this question by drawing upon the work of that most famous philosopher of categories, Aristotle. He provides examples of suggested upper-level ontologies which are currently in use, the Suggested Upper Merged Ontology (SUMO) and the Sowa Diamond, and argues that they are inferior to Aristotle’s upper-level categories. He then presents the upper-level category system Basic Formal Ontology (BFO), which was constructed under the influence of the Aristotelian table of categories, and makes the case for using BFO as the standard upper-level category system for reference ontologies.

Chapter 9 offers an example of the way in which Jansen’s considerations can be applied in one sort of theory that underpins the biomedical domain: the theory of the classification of living beings. On the basis of both philosophical and practical considerations, Heuer and Hennig justify the structure of the traditional, Linnaean, system of biological classification. Then they discuss certain formal principles governing the development of taxonomies in general, and show how classification in different domains must reflect the unique ontological aspects of the entities in each domain. They use these considerations to show that the traditional system of biological classification is also the most natural one, and thereby also the best.

Knowing how existing things are to be divided into categories is the first step in creating a reference ontology suitable for representing reality. But this is not enough. In addition to knowing what kinds of entities there are, we must know what kinds of *relations* they enter into with each other. We learn about the kinds of entities in reality by examining instances of these entities themselves. In ‘Ontological Relations’ (Chapter 10), Ulf Schwarz and Barry Smith argue that this is also the way to learn about the kinds of relations which obtain between these kinds of entities: we must examine the particular relations in which particular entities engage. They endorse the efforts of a group of leading ontological engineers, the Open Biomedical Ontologies (OBO) Consortium, to delineate the kinds of relations obtaining between the most general kinds of entities.

In Chapter 11, Ingvar Johansson offers a detailed treatment of one of the relations discussed in Chapter 10, the so-called *is_a* or subtype relation, which plays a particularly prominent role in information science. Johansson argues that there are good reasons to distinguish between four relations often confused when *is_a* relations are intended: genus-subsumption, determinable-subsumption, specification, and specialization. He shows that these relations behave differently in relation to definitions and so-called inheritance requirements. From the perspective predominant in this book, classifications should be marked by the feature of single inheritance: each species type in a classification should have a single parent-type or genus. The distinction between *single inheritance* and *multiple inheritance* is important both in information science ontologies and in some programming languages. Johansson argues that single inheritance is a good thing in subsumption hierarchies and is inevitable in pure specifications, but that multiple inheritance is often acceptable when *is_a* graphs are constructed to represent relations of specialization and in graphs that combine different kinds of *is_a* relations.

Many relations obtain between *continuant* entities; that is, entities, such as chairs and organisms, which maintain their identity through time. But reality also consists of processes in which continuant entities participate, which form a different category of entity, namely, *occurrent* entities. Just like continuants, occurrents can – and must – be classified by any information system which seeks a full representation of reality. For, just as there are continuants such as diseases, so there are the occurrents that are referred to in medicine as disease courses or disease histories. Hennig’s ‘Occurrents’ (Chapter 12) develops an ontology, or classification, of occurrent entities. He distinguishes between processes, which have what he calls an internal temporal structure, and other temporally extended occurrents, which do not. Further, he notes that certain important differences must be taken into account between types of occurrents and their instances. He argues that particular occurrents may instantiate more than one type at the same time, and that instances of certain occurrents are necessarily incomplete as long as they occur. By pointing out these and other important ways in which occurrents differ from continuants, Hennig’s work shows the urgency of the need for information systems to obtain clarity in their upper-level categories.

Finally, in Chapter 13, Johansson takes a wide-lens view of the junction of philosophy, ontology, and bioinformatics. He observes that some bioinformaticians, who work with terms and concepts, are reluctant to

believe that it is possible to have knowledge of mind-independent reality in the biological domain. He argues that there is no good reason for this tendency, and that it is even potentially harmful. For, at the end of the day, bioinformaticians cannot completely disregard the question as to whether the terms and concepts of their discipline refer to real entities. In the first part of the chapter, Johansson clarifies three different positions in the philosophy of science with which it would be fruitful for bioinformaticians to become familiar, defending one of them: Karl Popper's epistemological realism. In the second part, he discusses the distinction (necessary for epistemological realism) between the *use* and *mention* of terms and concepts, showing the importance of this distinction for bioinformatics.

This volume does not claim to have the final say in the new discipline of applied ontology. The main reason is that the ideas it presents are still being developed. Our hope is that we have made a case for the urgency of applying rigorous philosophical methods to the efforts of information scientists to represent reality. That urgency stems from the vast potential which such application can have for rendering information systems interoperable, efficient, and well-honed tools for the increasingly sophisticated needs of anyone whose life may be affected by scientific research – that is to say, of everyone. What the authors of this volume are working toward is a world in which information systems enable knowledge to be stored and represented in ways that do justice to the complexity of that information itself, and of the reality which it represents.

Acknowledgments

This book was written under the auspices of the Wolfgang Paul Program of the Humboldt Foundation, the European Union Network of Excellence on Medical Informatics and Semantic Data Mining, and the Volkswagen Foundation Project: Forms of Life. In addition, the authors would like to thank the following for valuable comments: Werner Ceusters, Pierluigi Miraglia, Fabian Neuhaus, Michael Pool, Steffen Schulze-Kremer, Cornelius Rosse, Dirk Siebert, Andrew Spear, and the participants of the First Workshop on Philosophy and Informatics in 2003. Thanks are due, also, to Michelle Carnell, Rachel Grenon, Robert Arp, and Dobin Choi.

Chapter 1: Philosophy and Biomedical Information Systems

Barry Smith and Bert Klagges

1. *The New Applied Ontology*

Recent years have seen the development of new applications of the ancient science of philosophy, and the new sub-branch of *applied philosophy*. A new level of interaction between philosophy and non-philosophical disciplines is being realized. Serious philosophical engagement, for example, with biomedical and bioethical issues increasingly requires a genuine familiarity with the relevant biological and medical facts. The simple presentation of philosophical theories and arguments is not a sufficient basis for future work in these areas. Philosophers working on questions of medical ethics and bioethics must not only familiarize themselves with the domains of biology and medicine, they must also find a way to integrate the content of these domains in their philosophical theories. It is in this context that we should understand the developments in *applied ontology* set forth in this volume.

Applied ontology is a branch of applied philosophy using philosophical ideas and methods from ontology in order to contribute to a more adequate presentation of the results of scientific research. The need for such a discipline has much to do with the increasing importance of computer and information science technology to research in the natural sciences (Smith, 2003, 155-166). As early as the 1970s, in the context of attempts at data integration, it was recognized that many different information systems had developed over the course of time. Each system developed its own principles of terminology and categorization which were often in conflict with those of other systems. It was for this reason that a discipline known as ontological engineering has arisen in the field of information science whose aim, ideally conceived, is to create a common basis of communication – a sort of Esperanto for databases – the goal of which would be to improve the compatibility and reusability of electronically stored information.

Various institutions have sprung up, including the Metaphysics Lab at Stanford University, the Ontology Research Group in Buffalo, New York, and the Laboratories for Applied Ontology in Trento, Italy. Research at these institutions is focused on the use of ontological ideas and methods in

the interaction between philosophy and various fields of information sciences. The results of this research have been incorporated into software applications produced by technology companies such as Ingenuity Systems (Mountain View, California), Cycorp, Inc. (Austin, Texas), and Ontology Works (Baltimore, Maryland). Rapid developments in information-based research technology have called forth an ontological perspective, especially in the field of biomedicine. This is illustrated in the work of research groups and institutions such as Medical Ontology Research at the US National Library of Medicine, the Berkeley Bioinformatics and Ontology Project at the Lawrence Livermore National Laboratory, the Cooperative Ontologies Programme of the University of Manchester, the Institute for Formal Ontology and Medical Information Science (IFOMIS) in Saarbrücken, Germany, and the Gene Ontology Consortium.

2. The Historical Background of Applied Ontology

The roots of applied ontology stretch back to Aristotle (384-322 BCE), and from the basic idea that it is possible to obtain philosophical understanding of aspects of reality which are at the same time objects of scientific research.

But how can this old idea be endowed with new life today? In order to answer this question, we must cast a quick glance back at the history of Western philosophy. An ontology can be seen, roughly, as a taxonomy of entities – objects, attributes, processes, and relations – in a given domain, complete with formal rules that govern the taxonomy (for a detailed exposition, see Chapter 2). An ontology divides a domain into *classes* or *kinds* (in the terminology of this volume, *universals*). Complex domains require multiple levels of hierarchically organized classes. Carl Linnaeus's taxonomies of organisms are examples of ontologies in this sense. Linnaeus also applied the Aristotelian methodology in medicine by creating hierarchical categories for the classification of diseases.

Aristotle himself believed that reality in its entirety could be represented with one single system of categories (see Chapter 8). Under the influence of René Descartes and Immanuel Kant, however, the focal point of philosophy shifted from (Aristotelian) *metaphysics* to *epistemology*. In a separate development, the Aristotelian-inspired view of categories, species, and genera as parts of a determined order came gradually to be undermined within biology by the Darwinian revolution. In the first half of the twentieth century, this two-pronged anti-ontological turn received

increasing impetus with the influence of the logical positivism of the Vienna Circle.

Toward the end of the twentieth century, however, there was another shift of ground, in philosophy as well as in biology. Philosophers such as Saul Kripke, Hilary Putnam, David Armstrong, Roderick Chisholm, David Lewis, and Ruth Millikan managed to bring ontological and metaphysical considerations back into the limelight of analytic philosophy under the title ‘analytical metaphysics’. This advance has brought elements of a still recognizably Aristotelian theory of categories (as the theory of universals or natural kinds) to renewed prominence. In addition, the growing importance of the new bioethics is helping to cast a new, ontological light on the philosophy of biology, above all in Germany in the work of Nikolaus Knoepffler and Ralf Stoecker.

In biology itself, traditional ideas about categorization which had been viewed as obsolete are now looked upon with favor once again. The growing significance of taxonomy and terminology in the context of current information-based biological research has created a terrain in which these ideas have blossomed once more. In fact, biology can be said to be enjoying a new golden age of classification.

3. *Ontological Perspectivalism*

One aspect of the Aristotelian view of reality still embraced by some ontologists is now commonly considered unacceptable, namely, that the whole of reality can be encompassed within one single system of categories. Instead, it is assumed that a multiplicity of ontologies – of partial category systems – is needed in order to encompass the various aspects of reality represented in diverse areas of scientific research. Each partial category system will divide its domain into classes, types, groupings, or kinds, in a manner analogous to the way in which Linnaeus’s taxonomies divided the domain of organisms into various upper-level categories (kingdom, phylum, class, species, and so forth), now codified in works such as the *International Code of Zoological Nomenclature* and the *International Code of Nomenclature of Bacteria*.

One and the same cross-section of reality can often be represented by various divisions which may overlap with one another. For example, the Periodic Table of the Elements is a division of (almost) all of material reality into its chemical components. In addition, the table of astronomical categories, a taxonomy of solar systems, planets, moons, asteroids, and so

forth, is a division of (the known) material reality – but from another perspective and at another level of granularity.

The thesis that there are multiple, equally valid and overlapping divisions of reality may be called *ontological perspectivalism* (see Chapter 6). In contrast to various perspectival positions in the history of Western philosophy – for example, those of Nietzsche or Foucault – this ontological variant of perspectivalism is completely compatible with the scientific view of the world. Ontological perspectivalism accepts that there are alternative views of reality, and that this same reality can be represented in different ways. The same section of the world can be observed through a telescope, with the naked eye, or through a microscope. Analogously, the objects of scientific research may be equally well-viewed or represented by means of a taxonomy, theory, or language.

However, the ontological perspectivalist is confronted with a difficult problem. How can these various perspectives be made compatible with one another? How can scientific disciplines communicate, and work together, if each treats of a different subdivision or granularity? Is there a discipline which can provide some platform for integration? In the following we will try to show that, in tackling this problem, there is no alternative to an ontology constructed from philosophically grounded, rigorous formal principles. Our task is practical in nature, and is subject to the same practical constraints faced in all scientific activity. Thus, even an ontology based on philosophical principles always will be a partial and imperfect edifice, which will be subject to correction and enhancement, so as to meet new scientific needs.

4. The Modular Structure of the Biological Domain

The perspectives relevant to our purposes in the domain of biomedical ontology are those which help us to formulate scientific explanations. These are often perspectives of a fine granularity, by means of which we gain insight into, for example, the number and order of genes on a chromosome, or the reactions within a chemical pathway. But if the scientific view of these structures is to have a significance for the goals of medicine, it must be seen through different, coarse-grained perspectives, including the perspective of everyday experience, which embraces entities such as diseases and their symptoms, human feelings and behavior, and the environments in which humans live and act.

As Gottfried Leibniz asserted in the seventeenth century, when perceived more closely than the naked eye allows, the entities of the natural world are revealed to be aggregates of smaller parts. For example, an embryo is composed of a hierarchical nesting of organs, cells, molecules, atoms, and subatomic parts. The ecological psychologist Roger Barker expresses it this way:

A unit in the middle range of a nesting structure is simultaneously both circumjacent and interjacent, both whole and part, both entity and environment. An organ – the liver, for example – is whole in relation to its own component pattern of cells, and is a part in relation to the circumjacent organism that it, with other organs, composes; it forms the environment of its cells, and is, itself, envired by the organism. (Barker, 1968, 154; compare Gibson, 1979)

Biological reality appears, in this way, as a complex hierarchy of nested levels. Molecules are parts of collections which we call cells, while cells are embedded, for example, in leaves, leaves in trees, trees in forests, and so forth. In the same way that our perceptions and behavior are more or less perfectly directed toward the level of our everyday experience, so too, the diverse biological sciences are directed toward various other levels within these complex hierarchies. There is, for example, not only clinical physiology, but also cell and molecular physiology; beside neuroanatomy there is also neurochemistry; and beside macroscopic anatomy with its sub-disciplines such as clinical, surgical, and radiological anatomy, there is also microscopic anatomy with sub-disciplines such as histology and cytology.

Ontological perspectivalism, then, should provide a synoptic framework in which the domains of these various disciplines can be linked, not only with each other, but also with an ontology of the granular level of the everyday objects and processes of our daily environment.

5. Communication among Perspectives

The central question is this: how do the coarse-grained parts and structures of reality, to which our direct perception and actions are targeted, relate to those finer-grained parts, dimensions, and structures of reality to which our scientific and technological capabilities provide access? This question recalls the project of the philosopher, Wilfrid Sellars, who sought what he called a stereoscopic view, the intent of which is to gather the content of our everyday thought and speech with the authoritative theories of the natural sciences into a single synoptic account of persons and the world

(Sellars, 1963). This stereoscopic view was intended to do justice, not only to the modern scientific image, but also to the manifest image of normal human reason, and to enable communication between them.

Which is the real sun? Is it that of the farmers or that of the astronomers? According to ontological perspectivalism, we need not decide in favor of the one or the other since both everyday and scientific knowledge stem from divisions which we can accept simultaneously, provided we are careful to observe their respective functions within thought and theory. The communicative framework which will enable us to navigate between these perspectives should provide a theoretical basis for treating one of the most important problems in current biomedicine. How do we integrate the knowledge that we have of objects and processes at the genetic (molecular) level of granularity with our knowledge of diseases and of individual human behavior, through to investigations of entire populations and societies?

Clearly, we cannot fully answer this question here. However, we will provide evidence that such a framework for integration can be developed as a result of the fact that biology and bioinformatics have implicitly come to accept certain theoretical and methodological presuppositions of philosophical ontology, presuppositions that pivot on an Aristotelian approach to hierarchical taxonomy.

Philosophical ideas about categories and taxonomies (and, as we will see, about many other traditional philosophical notions) have won a new relevance, especially for biology and bioinformatics. It seems that every branch of biology and medicine still uses taxonomic hierarchies as one foundation of its research. These include not only taxonomies of species and kinds of organisms and organs, but also of diseases, genomics and proteomics, cells and their components, biochemical reactions, and reaction pathways. These taxonomies are providing an indispensable instrument for new sorts of biological research in the form of massive databases such as Flybase, EMBL, Unigene, Swiss-Prot, SCOP, or the Protein Data Bank (PDB).¹ These allow new means of processing of data, resulting in extraction of information which can lead to new scientific results. Fruitful application of these new techniques requires, however, a solution to the problem of communication between these diverse category systems.

¹ See, for example, <http://www.cs.man.ac.uk/~stevensr/ontology.html>.

We believe that the new methods of applied ontology described in this volume bring us closer to a solution to this problem, and that it is possible to establish productive interdisciplinary work between biologists and information scientists wherein philosophers would act, in effect, as mediators.

6. *Ontology and Biomedicine*

There are many prominent examples of ways in which information technology can support biomedical research, including the coding of the human genome, studies of genetic expression, and better understanding of protein structures. In fact, all of these result from attempts to come to grips with the role of hereditary and environmental factors in health and the course of human diseases, and to search for material for new pharmaceuticals.

Current bioinformatics is extremely well-equipped to support calculation-intensive areas of biomedical research, focused on the level of the genome sequence, which can search for *quantitative* correlations, for example, through statistics-based methods for pattern recognition. However, an appropriate basis for *qualitative* research is less well-developed. In order to exploit the information we gain from quantitative correlations, we need to be able to process this information in such a way that we can identify those correlations which are of biological (and perhaps, clinical) significance. For this, however, we need a qualitative theory of types and relations of biological phenomena – an ontology – which also must include very general terms such as ‘object’, ‘species’, ‘part’, ‘whole’, ‘function’, ‘process’, and the like. Biologists have only a rather vague understanding of the meaning of these terms; but this suffices for their needs. Miscommunication between them is avoided simply in virtue of the fact that everyone knows which objects and processes in the laboratory are denoted by a given expression.

Information-technological processing requires explicit rigorous definitions. Such definitions can only be provided by an all-encompassing formal theory of the corresponding categories and relations. As noted already, information science has taken over the term ‘ontology’ to refer to such an all-encompassing theory. As is illustrated by the successes of the Gene Ontology (GO), developing such a resource can permit the mass of terminology and category systems thrown together in rather *ad hoc* ways over time to be unified within more overarching systems.

Already, the 1990s saw extensive efforts at modifying vocabularies in order to unite them within a common framework. Biomedical informatics offered framework approaches such as MeSH and SNOMED, as well as the creation of an overarching integration platform called the Unified Medical Language System (UMLS) (see National Library of Medicine). Little by little, the respective domains were indexed into robust and commonly accepted *controlled vocabularies*, and were annotated by experts to ensure the long-term compatibility and reusability of the electronically stored information. These controlled vocabularies contributed a great deal to the dawning of a new phase of terminological precision and orderliness in biomedical research, so that the integration of biological information that was hoped for seems achievable.

These efforts, however, were limited to the terminologies and the computer processes that worked with them. Much emphasis was placed upon the merely *syntactic* exactness of terms, that is, upon the grammatical rules applied to them as they are collected and ordered within structured systems. But too little attention was paid to the *semantic* clarity of these terms, that is, to their reference *in reality*. It was not that terms had no definitions – though such definitions, indeed, were often lacking. The problem was rather that these definitions had their origins in the medical dictionaries of an earlier time; they were written for people, not for computers. Because of this, they have an informal character, and are often circular and inconsistent. The vast majority of terminology systems today are still based on imprecisely formulated notions and unclear rules of classification.

When such terminologies are applied by people in possession of the requisite experience and knowledge, they deliver acceptable results. At the same time, they pose difficulties for the prospects of electronic data processing – or are simply inappropriate for this purpose. For this reason, the vast potential of information technology lies unexploited. For rigorously structured definitions are necessary conditions for consistent (and intelligent) navigation between different bodies of information by means of automated reasoning systems. While appropriately qualified, interested, and motivated people could make do with imprecisely expressed informational content, electronic information processing systems absolutely require exact and well-structured definitions (Smith, Köhler, Kumar, 2004, 79-94).

Collaboration between information scientists and biologists is all too often influenced by a variant of the *Star Trek* Prime Directive, namely,

‘Thou shalt not interfere with the internal affairs of other civilizations’. In the present context, these other civilizations are the various branches of biology, while ‘not to interfere’ means that most information scientists see themselves as being obliged to treat information prepared by biologists as something untouchable, and so develop applications which enable navigation through this information. Hence, information scientists and biologists often do not interact during the process of structuring their information, even though such interaction would improve the potential power of information resources tremendously. Matters are changing, now, with the development of OBI, the OBO Foundry Ontology for Biomedical Investigations (<http://obi.sourceforge.net/>), which is designed to support the consistent annotation of biomedical investigations, regardless of the particular field of study.

7. The Role of Philosophy

Up to now, not even biological or medical information scientists were able to achieve an ontologically well-founded means of integrating their data. Previous attempts, such as the Semantic Network of the UMLS (McCray, 2003, 80-84), brought ever more obvious problems stemming from the neglect of philosophical, logical, and especially definition-theoretical principles for the development of ontological theories to light (Smith, 2004, 73-84). Terms have been confused with concepts, while concepts have been confused with the things denoted by the words themselves and with the procedures by which we obtain knowledge about these things. Blood pressure has been identified, for example, with the *measuring* of blood pressure. Bodily systems, such as the circulatory system, have been classified as conceptual entities, but their parts (such as the heart) as physical entities. Further, basic philosophical distinctions have been ignored. For example, although the Gene Ontology has a taxonomy for functions and another for processes, initially there was no attempt to understand how these two categories relate or differ; both were equated in GO with ‘activity’. Recent GO documentation has improved matters considerably in these respects, with concomitant improvements in the quality of the ontology itself.

Since computer programs only communicate what has been explicitly programmed into them, communication between computer programs is more prone to certain kinds of mistakes than communication between people. People can read between the lines (so to speak), for example, by

drawing on contextual information to fill in gaps of meaning, whereas computers cannot. For this reason, computer-supported systems in biology and medicine are in dire need of maximal clarity and precision, particularly with respect to those most basic terms and relations used in all systems; for example, ‘*is_a*’, ‘*part_of*’, or ‘*located_in*’. An ontological theory based on logical and philosophical principles can, we believe, provide much of what is needed to supply this missing clarity and precision, and early evidence from the development of the OBO Foundry initiative is encouraging in this respect. This sort of ontological theory can not only support more coherent interpretations of the results delivered by computers, it also will enable better communication between, and among, the scientists of various disciplines. This is achieved by counteracting the fact that scientists bring a variety of different background assumptions to the table and, for this reason, often experience difficulties in communicating successfully.

One instrument for improving communication is the OBO Foundry’s Foundational Model of Anatomy (FMA) Ontology, developed through the Department of Biological Structure at the University of Washington in Seattle, which is a standard-setter among bioinformation systems. The FMA represents the structural composition of the human body from the macromolecular level to the macroscopic level, and provides a robust and consistent schema for the classification of anatomical unities based upon explicit definitions. This schema also provides the basis for the Digital Anatomist, a computer-supported visualization of the human body, and provides a pattern for future systems to enable the exact representation of pathology, physiological functions, and the genotype-phenotype relations.

The anatomical information provided by the FMA Ontology is explicitly based upon Aristotelian ideas about the correct structure of definitions (Michael, Mejino, Rosse, 2001, 463-467). Thus, the definition of a given class in the FMA – for example, the definition for ‘heart’ or ‘organ’ – specifies what the corresponding instances have in common. It does this by specifying (a) a *genus*, that is, a class which encompasses the class being defined, together with (b) the *differentiae* which characterize these instances within the wider class and distinguish them from its other members. This modular structure of definitions in the FMA Ontology facilitates the processing of information and checking for mistakes, as well as the consistent expansion of the system as a whole. This modular structure also guarantees that the classes of the ontology form a genuine categorial tree in the ancient Aristotelian sense, as well as in the sense of the Linnaean taxonomy. The Aristotelian doctrine, according to which

definition occurs via the nearest genus and specific difference, is applied in this way to current biological knowledge.

In earlier times the question of which types or classes are to be included within the domain of scientific anatomy was answered on the basis of visual inspection. Today, this question is the object of empirical research within genetics, along with a series of related questions concerning, for example, the evolutionary predecessors of anatomical structures extant in organisms. In course of time, a phenomenologically recognizable anatomical structure is accepted as an instance of a genuine class by the FMA Ontology only after sufficient evidence is garnered for the existence of a structural gene.

8. *The Variety of Life Forms*

The ever more rapid advance in biological research brings with it a new understanding of the variety of characteristics exhibited by the most basic phenomena of life. On the one hand, there is a multiplicity of *substantial forms of life*, such as mitochondria, cells, organs, organ systems, single- and many-celled organisms, kinds, families, societies, populations, as well as embryos and other forms of life at various phases of development. On the other hand, there are certain basic building blocks of *processes*, what we might call *forms of processual life*, such as circulation, defence against pathogens, prenatal development, childhood, adolescence, aging, eating, growth, perception, reproduction, walking, dying, acting, communicating, learning, teaching, and the various types of social behavior. Finally, there are certain types of processes, such as cell division or the transport of molecules between cells, in every phase of biological development.

Developing a consistent system of ontological categories founded upon robust principles which can make these various forms of life, as well as the relations which link them, intelligible requires addressing several issues which are often ignored in biomedical information systems, or addressed in an unsatisfactory manner, because they are philosophical in nature. These issues show the unexplored practical relevance of philosophical research at the frontier between information science and empirical biology.² These issues include:

² See also: Smith, Williams, Schulze-Kremer, 2003, 609-613; Smith, Rosse, 2004, 444-448.

(1) Issues pertaining to the different modes of existence through time of diverse forms of life. Substances (for example, cells and organisms) are fundamentally different from processes with respect to their mode of existence in time. Substances exist as a whole at every point of their existence; they maintain their identity over time, which is itself of central relevance to the definition of 'life'. By contrast, processes exist in their temporal parts; they unfold over the course of time and are never existent as a whole at one and the same instant (Johansson, 1989; Grenon, Smith, 2004, 69-103).

We can distinguish between entities which exist continually (continuants) and entities which occur over time (occurrents). It is not only substances which exist continually, but also their states, dispositions, functions, and qualities. All of these latter entities stand in certain relations on the one hand to their substantial bearers and on the other hand to certain processes. For example, functions are generally *realized* in processes. In the same way that an organism has a *life*, a *disposition* has the possibility of being *realized*, and a *state* (such as a disease) has its *course* or its *history* (which can be represented in a medical record).

(2) The notion of *function* in biology also requires analysis. It is not only genes which have functions that are important for the life of an organism; so do organs and organ systems, as well as cells and cellular parts such as mitochondria or chloroplasts. A function inheres in a body part or trait of an organism and is realized in a process of *functioning*; hence, for example, one function of the heart is to pump blood. But what does the word 'function' mean in this context? Natural scientists and philosophers of science from the twentieth century have deliberately avoided talk of functions – and of any sort of teleology – because teleological theories were seen to be in disagreement with the contemporary scientific understanding of causation. Yet, functions are crucial for the worldview (the ontology) of physicians and medical researchers, as a complete account of a body part or trait often requires reference to a function. Further, it is in virtue of the body's ability to transform malfunctioning into functioning that life persists.

The nature of functions has been given extensive treatment in recent philosophy of biology. Ruth Millikan, for example, has offered a theory of proper function as a disposition belonging to an entity of a certain type, which developed over the course of evolution and is responsible for (at least in part) the existence of more entities of its type (Millikan, 1988). However, an entity has a function only within the context of a biological

system and this requires, of course, an analysis of *system*. But existing philosophical theories lack the requisite precision and general application necessary for a complete account of functions and systems (Smith, Papakin, Munn, 2004, 39-63; Johansson, *et al.*, 2005, 153-166).

(3) The issue of the components and structure of organisms also needs to be addressed. In what relation does an organism stand to its body parts? This question is a reappearance of the ancient problem of *form* and *matter* in the guise of the problem of the relation between the organism as an organized whole, and its various material bearers (nucleotides, proteins, lipids, sugars, and so forth). Single-celled as well as multi-celled organisms exhibit a certain modular structure, so that various parts of the organism may be identified at different granular levels. There are a variety of possible partitions through which an organism and its parts can be viewed depending upon whether one's focus is centered on molecular or cellular structures, tissues, organ systems, or complete organisms. Because an organism is more than the sum of its parts, this plurality of trans-granular perspectives is central to our understanding of an organism and its parts. The explanation of how these entities relate to one another from one granular level to the next is often discussed in the literature on emergence, but is seldom imbued with the sort of clarity needed for the purposes of automated information representation.

The temporal dimension contains modularity and corresponding levels of granularity as well. So, if we focus successively on seconds, years, or millennia, we perceive the various partitions of processual forms of life, such as individual chemical reactions, biochemical reaction paths, and the life cycles of individual organisms, generations, or evolutionary epochs.

(4) We also need to address the issue of the nature of biological kinds (species, types, universals). Any self-respecting theory of such entities must allow room for the *evolution of kinds*. Most current approaches to such a theory appeal to mathematical set theory, with more or less rigor. A biological kind, however, is by no means the same as the set of its instances. For, while the identity of a set is dependent upon its elements or members and, hence, participates to some degree in the world of time and change, sets themselves exist outside of time. By contrast, biological kinds exist in time, and they continue to exist even when the entirety of their instances changes. Thus, biological kinds have certain attributes in common with individuals (Hull, 1976, 174-191; Ghiselin, 1997), and this is an aspect of their ontology which has been given too little attention in bioinformatics.

Existing bioinformation systems concentrate on terms which are organized into highly general taxonomical hierarchies and, thus, deal with biological reality only at the level of classes (kinds, universals). Individual organisms – which are instantiations of the classes represented in these hierarchies – are not taken into consideration. This lack of consideration has partially to do with the fact that the medical terminology, which constitutes the basis for current biomedical ontologies, so overwhelmingly derives from the medical dictionaries of the past. Authors of dictionaries, as well as those involved in knowledge representation, are mainly interested in what is general. However, an adequate ontology of the biological domain must take individuals (instances, particulars) as well as classes into account (see Chapters 7, 8, and 10). It must, for example, do justice to the fact that biological kinds are always such as to manifest, not only typical instances, but also a penumbra of borderline cases whose existence sustains biological evolution. As we will show in what follows, if we want to avoid certain difficulties encountered by previous knowledge representation systems, the role of instances in the structuring of the biological domain cannot be ignored.

(5) There is much need, also, for a better understanding of synchronic and diachronic identity. Synchronic identity has to do with the question of whether x is the same individual (protein, gene, kind, or organism) as y , while diachronic identity concerns the question of whether x is *today* the same individual (protein, gene, kind, or organism) as x was yesterday or a thousand years ago. An important point of orientation on this topic is the logical analysis of various notions of identity put forward by the *Gestalt*-psychologist Kurt Lewin (Lewin, 1922). Lewin distinguishes between physical, biological, and evolution-theoretic identity; that is, between the modes of temporal persistence of a complex of molecules, of an organism, or of a kind. Contemporary analytic philosophers, such as Eric Olson or Jack Wilson, have also managed to treat old questions (such as those of personal identity and individuation) with new ontological precision (Olson, 1999; Wilson, 1999).

(6) There is also a need for a theory of the role of environments in biological systems. Genes exist and are realized only in very specific molecular contexts or environments, and their concrete expression is dependent upon the nature of these contexts. Analogously, organisms live in niches or environments particular to them, and their respective environments are a large part of what determine their continued existence.

However, the philosophical literature since Aristotle has shed little light upon questions relating to the ontology of the environment, generally according much greater significance to substances and their accidents (qualities, properties) than to the environments surrounding these substances. But what *are* niches or environments, and how are the dependence relations between organisms and their environments to be understood ontologically? The relevance of these questions lies not only within the field of developmental biology, but also ecology and environmental ethics, and is now being addressed by the OBO Foundry's new Environment Ontology (<http://environmentontology.org>).

9. *The Gene Ontology*

The rest of this volume will provide examples of the methods we are advocating for bringing clarity to the use of terms by biologists and by bioinformation systems. We will conclude this chapter with a discussion of the Gene Ontology (see Gene Ontology Consortium, ND), an automated taxonomical representation of the domains of genetics and molecular biology. Developed by biologists, the Gene Ontology (GO) is one of the best known and most comprehensive systems for representing information in the biological domain. It is now crucial for the continuing success of endeavors such as the Human Genome Project, which require extensive collaboration between biochemistry and genetics. Because of the huge volumes of data involved, such collaboration must be heavily supported by automated data exchange, and for this the controlled vocabulary provided by the GO has proved to be of vital importance.

By using humanly understandable terms as keys to link together highly divergent datasets, the GO is making a groundbreaking contribution to the integration of biological information, and its methodology is gradually being extended, through the OBO Foundry, to areas such as cross-species anatomy and infectious disease ontology.

The GO was conceived in 1998, and the Open Biomedical Ontologies Consortium (see OBO, ND) created in 2003, as an umbrella organization dedicated to the standardization and further development of ontologies on the basis of the GO's methodology. The GO includes three controlled vocabularies – namely, *cellular component*, *biological process*, and *molecular function* – comprising, in all, more than 20,000 biological terms. The GO is not itself an integration of databases, but rather a vocabulary of terms to be used in describing genes and gene products. Many powerful

tools for searching within the GO vocabulary and manipulation of GO-annotated data, such as AmiGO, QuickGO, GOAT, and GoPubMed (see GOAT, 2003 and gopubmed.org, 2007), have been made available. These tools help in the retrieval of information concerning genes and gene products annotated with GO terms that is not only relevant for theoretical understanding of biological processes, but also for clinical medicine and pharmacology.

The underlying idea is that the GO's terms and definitions should depend upon reference to individual species as little as possible. Its focus lies, particularly, on those biological categories – such as *cell*, *replication*, or *death* – which reappear in organisms of all types and in all phases of evolution. It is not a trivial accomplishment on the GO's part to have created a vocabulary for representing such high-level categories of the biological realm, and its success sustains our thesis that certain elements of a philosophical methodology, like the one present in the work of Aristotle, can be of practical importance in the natural sciences.

Initially, the GO was poorly structured and some of its most basic terms were not clearly defined, resulting in errors in the ontology itself. (See: Smith, Köhler, Kumar, 79-94; Smith, Williams, Schulze-Kremer, 609-613). The hierarchical organization of GO's three vocabularies was similarly marked by problematic inconsistencies, principally because the *is_a* and *part_of* relations used to define the architecture of these ontologies were not clearly defined (see Chapter 11).

In early versions of the GO, for example, the assertions such as 'cell component *part_of* Gene Ontology' existed alongside properly ontological assertions such as 'nucleolus *part_of* nuclear lumen' and 'nuclear lumen *is_a* cellular component'. Unlike the second and third assertions, which rightly relate to part-whole relations on the side of biological reality, the first assertion captures an inclusion relation between a term and a list of terms in the GO itself. This misuse of '*part_of*' represents a classic confusion of *use* and *mention*. A term is *used* if its meaning contributes to the meaning of the including sentence, and it is merely *mentioned* if it is referred to, say in quotation marks, without taking into account its meaning (for more on this distinction and its implications, see Chapter 13).

10. Conclusion

The level of philosophical sophistication among the developers of biomedical ontologies is increasing, and the characteristic errors by which

such ontologies were marked is decreasing as a consequence. Major initiatives, such as the OBO Foundry, are a reflection of this development, and further aspects of this development are outlined in the chapters which follow.

Chapter 2: What is Formal Ontology?

Boris Hennig

1. *Ontology and Its Name*

‘Ontology’ is a neologism coined in early modern times from Greek roots. Its meaning is easy to grasp; *on* is the present participle of the Greek *einai*, which means ‘to be’, and *logos* derives from *legein*, ‘to talk about’ or ‘to give an account of’ something. Accordingly, ontology is the discourse that has being as its subject matter. This is what Aristotle describes as *first philosophy*, ‘a discipline which studies that which is, insofar as it is, and those features that it has in its own right’ (*Meta.* Γ1, 1003a21-2).³

In a sense, every philosophical or scientific discipline studies things that exist. Yet, the term ‘ontology’ does not apply to every discipline that studies that which is. Although sciences do deal with features of existing things, they do not deal with them insofar as they exist. Special sciences study only certain kinds of things that exist, and only insofar as these things exhibit certain special features. Two different kinds of restrictions are involved in circumscribing what a special science is. A special science either studies only a limited range of things, or it studies a limited aspect of the things it studies. Physics, for instance, studies the physical properties of everything that has such properties. Biology only studies living beings and only insofar as they are alive, not insofar as they are sheer physical objects. Differential psychology studies human beings insofar as they differ from other human beings in ways that are psychologically measurable. Further, two different special sciences may very well have overlapping domains, that is, domains that include the same members. For example, the claims of physics and chemistry apply to the very same things, except that the former investigates their physical properties, while the latter their chemical properties.

Ontology differs from such sciences as physics and differential psychology, but not because it considers another special range of things. Every object studied by ontology is also studied by some other discipline. However, ontology studies a different *aspect* of those things. According to Aristotle, ontology is concerned with everything that exists only *insofar as it exists*. Existence itself is the aspect relevant to ontology. Hence, ontology will be possible only if there are features that each existing thing has only

³ All translations are the author’s unless otherwise specified.

because, and insofar as, it exists. Momentarily, we will ask what sorts of features these may be. The objective of this section, however, is to give a preliminary impression of what ontology is by considering the history of the discipline and its name.

Although Aristotle's *Metaphysics* already deals with questions of ontology, the word 'ontology' is much younger than this work. As a title for a philosophical discipline, *ontologia* has been in use since about the seventeenth century. Jacob Lorhard, rector of a German secondary school, uses this term in his *Ogdoas Scholastica* (1606) as an alternative title for metaphysics as it was taught in his school.⁴ However, he does not explain the term further. The book does not contain much more than a set of tree diagrams with the root node of one of them labelled, *metaphysica seu ontologia*. More prominently, the German philosopher Christian Wolff uses 'ontologia' in 1736 as a name for the discipline introduced by Aristotle in the passage quoted above (Wolff, 1736). The list of topics that Wolff discusses under this heading resembles the one given by Lorhard. It includes the notion of being, the categories of quantity and quality, the possible and the impossible, necessity and contingency, truth and falsehood, and the several kinds of causes distinguished in Aristotelian physics (material, efficient, formal, and final). This choice of topics certainly derives from Aristotle's *Metaphysics* and such works as the *Metaphysical Disputations* (1597) by Francisco Suárez.

We can gather some additional facts about the early use of the term 'ontologia' by considering the first known appearance of the corresponding adjective in the *Lexicon Philosophicum* (1613) by Rudolph Goclenius. A foray into his use of 'ontological' will provide insight into how the term came to be used as it today; but, as we will see, there are some important respects in which his usage differs from contemporary usage (and, thus, from the usage in this volume). Goclenius uses 'ontological' in his entry on *abstraction*, where he discusses *abstraction of matter*. As everywhere else in his lexicon, he does not present a unified account of the phenomenon in question, but rather lists several definitions and other findings from the literature. In the present context, we are not concerned with what Goclenius means by *abstraction* and *matter*, although the concept of matter will become important later in our discussion of formal ontology. Provisionally, matter can be taken to be the stuff out of which a thing is made. To abstract it from a thing simply means to take it away from that

⁴ The second edition appeared in 1613 under the title *Theatrum Philosophicum*.

thing, in our imagination or in reality. For the time being, we are primarily interested in the sense in which Goclenius uses the epithet ‘ontological’. In science, he says, there are three different ways of abstracting matter from given things.

First, one may ignore the particular lump of matter out of which a given thing is made, but still conceive of the thing as being made up of some matter or other. According to Goclenius, this is what natural scientists do: they investigate particular samples, and they study their material nature. They are only interested in one sample, rather than another, when the samples differ with respect to their general properties. In studying a particular diamond, for instance, scientists ignore its particularity and consider only those features that any other diamond would have as well. Scientists abstract from a particular thing’s matter in order to grasp those general features of a thing in virtue of which it falls under a certain category; but the fact that things of its type are made of some matter or other remains a factor in their account. This is what Goclenius calls *physical* abstraction.

Second, we may ignore all matter whatsoever, in such a way that no matter at all figures in our account of the subject under investigation. This kind of abstraction is practiced in geometry and, accordingly, Goclenius calls it *mathematical* abstraction. But he also calls it *ontological* abstraction, glossing the latter term as ‘pertaining to the philosophy of being and of the transcendental attributes’ (Goclenius, 1613, 16). We will explain this phrase in due course.

Finally, Goclenius continues, one may abstract matter from a given thing in reality as much as in thought. The result will be that the entity in question literally no longer possesses any matter. This Goclenius calls *transnatural* abstraction, of which, he claims, only God and the so-called divine Intelligences are capable.

There are at least three important things to note here. First, Goclenius identifies ontological abstraction with mathematical abstraction. He thereby implies that ontology in general, as much as mathematics, is concerned with abstract entities and formal structures. For instance, geometry is concerned with the properties that physical objects have only by virtue of their shape and location. Their other properties, such as color, weight, smell, etc., are irrelevant. In this sense, geometry abstracts from the matter that is shaped and focuses on the shapes themselves. Whether a triangle is made of iron or wood makes no geometrical difference. If formal ontology is abstract in the same sense, it should also abstract from

certain properties of things and focus on their more general features. Later, we will explore what these more general features might be. What is important here is that ontological abstraction goes farther than mere physical abstraction. The physicist is not interested in particular samples but, rather, in material things insofar as they are material. According to Goclenius, ontology is not interested in matter at all; since concrete things are composed of matter, ontology is not concerned with concrete things at all, not even in a general way.

Second, Goclenius equates ontology with the philosophical doctrine of the transcendental attributes. These attributes include *being*, *oneness* (or *unity*), *goodness*, and *truth*. *Being* and *oneness* are discussed by Aristotle; *goodness* and *truth* are introduced by later authors (Aertsen *et al.*, 1998). These attributes are called ‘transcendental’ because they apply to every existing thing, regardless of any categorial boundary. That is, they surpass (or transcend) the general categories which distinguish things of different kinds. According to neo-Platonic authors like Pseudo-Dionysius the Areopagite, the transcendental notions not only surpass the categorial boundaries between things, but also the realm of the things to which they apply, that is, the entire world. For instance, Dionysius writes that the ‘name being extends to all beings which are; and it is beyond them’ (Pseudo-Dionysius, *The Divine Names*, 5, 816B). In other words, the transcendental attributes are manifested by everything in the world, but they do not apply to themselves. The transcendental attribute *being* is not itself something that exists. In any case, if ontology studies the features of everything that exists insofar as it exists, then it will also be concerned with the transcendental attributes.

Third, Goclenius does not use the epithet ‘ontological’ in order to indicate something that *really* or *actually happens*. When we ontologically abstract matter from a thing, we do not really take away its matter. We do so only in thought. Real abstraction, by contrast, is what Goclenius calls *transnatural* abstraction, and it occurs when God separates the human soul from its body. In this regard, Goclenius’ use of ‘ontological’ is directly opposed to some of the contemporary uses of this word. When contemporary writers call something ontological, they often mean to indicate that it really obtains, or at least that it has implications for what exists independently of our thoughts. Thus, ontology is often opposed to epistemology; the former is often said to be about what there really is, whereas the latter is only about what we know. A common view, for instance, is that ontology is concerned with the level of things rather than

the level of truths. That is, ontology concerns objects in the world, not our thoughts as they are expressed in true propositions or true sentences (Smith and Mulligan, 1983, 73). Goclenius, however, does not distinguish between things and truths. For him, both *being* and *truth* are transcendental attributes that apply to everything that is, on every conceivable level. Accordingly, he has no reason to suppose that it makes a difference whether we study things or truths, and ontology may be the study of both.

In fact, Goclenius' use of the epithet 'ontological' differs from the modern one in all three respects that we have emphasized. First, ontology is no longer considered to be as abstract as mathematics. It does not abstract from all matter whatsoever, since it must also discuss the general features that things possess by virtue of being material and particular. Second, at least in the tradition of analytic philosophy, ontology does not include a treatment of such transcendental attributes as goodness and truth. Instead, these topics are dealt with in ethics and epistemology. Third, as we have seen, the opposition of real and ontological abstraction appears odd from a modern perspective.

We will see, however, that there is also some continuity between Goclenius' and contemporary uses of 'ontology'; ontology is still considered an abstract discipline in the sense that it avoids dependence on particular references. Further, the idea that there are at least two transcendental attributes which surpass the categorial boundaries – namely, oneness and being – is still upheld. Finally, many contemporary thinkers certainly would contrast ontological features and happenings with *transnatural* ones, that is, features and happenings that surpass the realm of nature.

2. *Some Things that are not Formal Ontology*

So far, we have introduced a rough notion of ontology as the study of features that things have insofar as they exist, and not insofar as they are concrete objects consisting of this rather than that matter. Since ontology, conceived in this way, abstracts from matter in the same way in which mathematics abstracts from matter, ontology would seem to be *formal ontology*.

What is formal ontology? Edmund Husserl, who introduced this term into philosophy, describes it as the 'eidetic science of the object as such' (Husserl, *Husserliana*, 3/1, 26-27). *Eidetic* derives from the Greek *eidos*, which means *form*. Therefore, we will approach Husserl's formula by

means of a brief discussion of the general distinction between matter and form. This will lead to a discussion of experience and its objects, thus enabling us to understand the second part of Husserl's description of the object as such.

Husserl deliberately uses the term 'eidetic' instead of 'formal', because he wants to avoid misleading connotations (Husserl, *Husserliana*, 3/1, 9). He is well-advised in doing so, since there are at least two common – and mistaken – accounts of what it means for a discipline to be formal.

First, a discipline is sometimes called formal merely because its claims are expressed by means of formal symbolism or even only a shorthand notation, as when one writes ' $\forall x: \text{MAN}(x) \rightarrow \text{MORTAL}(x)$ ' instead of 'all men are mortal'. Shorthand notations, however, are merely short, and sometimes not even that. There is no particular reason for calling them formal. Logic and mathematics are indeed formal disciplines, and they often use shorthand notations. But logic and mathematics are not formal *because* they use this kind of symbolism. For one thing, mathematical and logical truths can be expressed perfectly well in prose, although this would often take up more space. For another, any old body of knowledge can be expressed by short and rigorously defined symbols without, thereby, turning into a formal discipline. Logic and mathematics are properly called 'formal' only because they are *about* formal structures and features; for instance, those of shorthand symbolisms. Hence, formal ontology may indeed use symbolic shorthand notations as far as they are helpful; but it need not do so, and it will not be formal by virtue of doing so.

Second, formal ontology has sometimes been opposed to *regional* or *material* ontology, and both labels – 'formal' and 'material' – were introduced by Husserl (Husserl, *Ideen*, §9). There are separate regional ontologies for the domains of physics, biology, differential psychology, and so forth. It has been claimed that formal disciplines are 'set apart from regional or material disciplines in that they apply to all domains of objects whatsoever, so that they are independent of the peculiarities of any given field of knowledge' (Smith and Smith, 1995, 28). According to this view, formal ontologists should only advance judgments that hold true of all objects in general. This is not far from the truth, but some qualifications are in order. For example, it is not the case that every claim that is made within formal ontology applies to everything that exists. Formal ontology can also study the formal features of a limited range of entities, in the same way in which geometry can study the shapes of a limited range of entities.

Admittedly, it is difficult to say what it means for a discipline or judgment to be *about* or *apply to* something. For instance, it is not clear whether ‘beavers are rodents’ is about beavers, about rodents, or about the whole world. For in some sense, all judgments are about and apply to the entire world and everything in it. It holds true of the world that, in it, beavers are rodents. We will make the simplistic assumption that judgments apply to the things that are explicitly mentioned in them. On this basis, formal disciplines explicitly mention everything that exists by using very general and abstract descriptions, whereas regional and material disciplines mention only some of existing things, but presumably in more detail. Thus, whether a discipline is formal or not depends on the entities to which its claims refer, and on the way in which it refers to them.

However, there are two quite different ways in which a judgment may be said to explicitly mention or refer to particular objects.

(1) Judgments like ‘Marlene Dietrich was beautiful’ or ‘that child over there is intelligent’ are *particular* judgments.⁵ Particular *objects* are concrete, discrete, and they exist only once. Particular *judgments* refer to such things by using proper names or demonstrative expressions like ‘Marlene’, ‘this’, or ‘over there’. Further, their truth depends on the state of exactly those particular things to which they refer.

(2) The other way in which a judgment refers to specific things in the world consists in its being *specific*. Specific judgments hold true only of a limited range of entities, such as the judgment ‘some actresses are beautiful’. This judgment holds true only of actresses, and not of other persons or things. Although the truth of specific judgments still depends on the state of particular things, they do not refer to these things by using a demonstrative or name. They are, as it were, about *anonymous* particular objects. Specific judgments do not apply to everything in general; but they refer to their objects by means of a general form which may single out an unspecified number of particular objects.

This distinction between particular and specific judgments is important because it will turn out that a formal ontological theory may only advance specific judgments, but not particular ones. Thus, the point is not that formal ontology applies to all objects alike, but rather that it applies to certain ranges of objects that may be referred to by means of general terms. Whereas formal ontology must not refer to *particular* beings like Marlene Dietrich or that child over there, it can still refer to *specific* kinds of beings

⁵ The distinction is also drawn by Kant, 1781, B95. However, translations usually use ‘singular’ where we use ‘particular’, and ‘particular’ where we use ‘specific’.

like organisms in general or anonymous children and actresses. Hence, formal ontology may indeed advance judgments about the specific entities within a limited domain of knowledge, as long as none of these judgments are *particular* ones. Ontology is formal as long as it picks out and applies to particular entities solely by referring to general aspects of them; in other words, to some aspect of their *forms*.

Thus, formal ontology is not the same as general ontology (which would deal only with features that all things share) and, hence, it is not opposed to material or regional ontology. Rather, an ontological theory may be formal and regional at the same time. A regional ontology deals with a limited range of entities, but as long as it does not advance any particular judgments, it can still deal with them in a general way. For instance, the (regional) ontology of occurrents found in Basic Formal Ontology (see Chapter 1) studies only a limited range of entities, namely those that occur or unfold in time, but it does not study specific events or processes in particular, such as the death of Socrates or the Great Depression.

But there are still several sciences, such as physics and chemistry, which study specific phenomena in a general way. We have not yet found a way of distinguishing them from formal ontology.

3. *Matter and Form According to Aristotle*

General is the opposite of *regional*, and *formal* is the opposite of *material*. Formal ontology, rather than being non-regional, is non-*material*. It may study a specific kind of thing, but that does not mean that it studies particular and concrete instances of these kinds. What does this mean? In order to distinguish formal from material ontology, we will now consider the distinction between matter and form. There are at least two different traditional conceptions of the difference between matter and form, which are attributable to Aristotle and Immanuel Kant respectively. This section discusses two ways of drawing the distinction which we will call Aristotelian. We will turn to Kant in the next section.

Aristotle develops the distinction between form and matter in his treatment of movement and change. In his *Physics*, he characterizes matter as the primary underlying substrate from which a concrete thing comes into being and which persists in this thing (II, 3). This might be taken to mean that matter is the persisting subject of any kind of change. But this definition is not tenable, since an organism may change with regard to its

matter; it may, for example, gain and lose parts, yet remain the same organism. In this case the organism, not its matter, is the persisting subject of change. Hence, not everything that underlies and persists during a change can be matter. Presumably, what Aristotle meant is that it makes sense to speak of matter only in contexts where some change is possible. The unchanging does not consist of matter, but not everything that may change is, thereby, matter. The result, then, is that he does not provide a complete account of what matter is.

Within the later Aristotelian tradition, matter is often identified with the *principle of individuation* of material things.⁶ This means that the matter of a thing is what makes it *this* rather than *that* thing. Even when things have the same properties and, hence, bear the exact same form, they can differ from each other merely by being made up of different parcels of matter. This brings us back to the above remarks on particularity. Concrete things are particular in virtue of the fact that they are made of matter. To be particular is to exist only once, at some unique location in time and space, and this is why we can refer to particular things in their particularity by using demonstrative expressions like ‘here’ and ‘now’. Therefore, to be material can, in most cases, be taken to be coextensive with being subject to possible reference by demonstratives.⁷ When a concrete thing is referred to by means of a demonstrative, it is not specified in terms of its general *form*, but in terms of its matter. Accordingly, we may claim that demonstratives introduce elements into discourse that are non-formal, that is, material. That formal ontology must not refer to matter will then mean that it must not advance judgments that contain demonstrative expressions (Cf. Husserl, *Ideen*, §7, *Husserliana* 3/1, 21).

But why should formal ontology not employ demonstrative expressions? To be here rather than there, or to occur now rather than earlier, certainly is a formal feature of a thing that it can share with other things. Further, geometry is concerned with exactly such features that objects have by virtue of being here or there, or extending from here to there; and geometry is certainly a formal discipline. The things in our world are *in general* here, now, there, or then. Therefore, any useful formal ontology that applies to real objects should also include a treatment of

⁶ See, for instance, Aquinas’ commentary on Boethius’ *De Trinitate*, II, q. 4, a. 1-2, *Opera* (editio Leonina) vol. 50. Cf. Charlton, 1972.

⁷ This only holds in most cases. Points in space are particular without being material, and there may be immaterial, but particular, things. Aristotle sometimes speaks of intelligible matter in such cases (e.g., *Metaphysics* Z 10).

space and time. But in order to point out the merely spatial difference between two locations, we must employ demonstrative expressions: one of them is ‘here’ (or at this and that location relative to here), another one is ‘over there’. On the face of it, then, the second Aristotelian conception, according to which matter is the principle of individuation of concrete things, is also of no help when it comes to circumscribing formal ontology. It seems that formal ontology must employ demonstrative expressions after all.

However, we will see that this is not the case. Formal ontology must not refer to objects by means of judgments containing demonstrative expressions. But how is this possible, given that formal ontology must include references to space and time? In what follows, we will see that there is a difference between an ontology that *uses* demonstrative reference in order to identify its objects, and an ontology that *reflects upon the use* of demonstratives, but without using them. This distinction is attributable to Kant. We will now explain it in more detail by turning to Kant’s conception of the contrast between matter and form, which differs from the Aristotelian one in several important respects.

4. *Kant on Formal Content*

We are still looking for an understanding of ‘form’ that enables us to grasp the distinction between formal and material ontology. Kant writes that the concepts of matter and form are ‘concepts of reflection’. This means that they are properly used in reflective judgments. In Kant’s own terms, reflective judgments express the ‘consciousness of the relation of given representations to the different sources or faculties of cognition’ (Kant, 1781, B316). In more familiar terms, they state how sense impressions, perceptions, and cognitions relate to the faculties that make them possible. If *form* and *matter* are concepts of reflection, they are concepts that figure prominently, or even exclusively, in judgments about how our sense impressions, perceptions, and cognitions relate to the faculties that make them possible.

According to Kant, we may achieve knowledge about the world by combining two sources of cognition. The first of these sources Kant calls *intuition*, which is the capacity or act of representing concrete particular objects, whether real or imaginary, to the mind. But intuition alone does not suffice for cognition. ‘Thoughts without content are void’, Kant claims, and ‘intuitions without concepts are blind’ (Kant, 1781, B75). Hence,

intuition must be supplemented by what Kant calls *judgment*, the act or capacity of uniting representations under concepts such as ‘existence’, ‘unity’, ‘substance’, or ‘cause’. By bringing representations under such concepts, understanding turns a subjective representation into an objective experience of a real object. For example, a mere sequence of visual experiences is brought under the concept of causation; it turns into an experience of a causal process. Kant claims that all cognition of empirical objects must work in this manner; thus, he is not only talking about our human cognitive abilities, but about what it would take for any rational being to experience an empirical object.

We are interested in the way in which Kant draws the distinction between two sorts of content that an experience may have, namely, the *material content* and the *formal content*. The distinction between the two sources of cognition, i.e., intuition and understanding, does not straightforwardly map onto this distinction. Put differently, material content is not quite the same as what Kant calls *empirical* content. The *empirical* content of experience is supplied by sense perception or other kinds of intuition and, thus, consists in the representation of particular concrete objects. However it is important to see that, according to Kant, the *formal* content of experience is not, in turn, exclusively supplied by our understanding (which means that not all empirical content is material). Rather, when intuition provides us with the representations of concrete things and locations in space and time, it has already introduced its own forms. According to Kant, the pure forms of intuition are space and time.

Kant’s distinction between the formal and material content of experience can be understood in the following way. In order to achieve knowledge about any given thing, we must first establish a relation to that thing. We need to relate to it by means of some of its properties, by looking at it, by pointing to it, or by using its proper name. For instance, in order to find out how beavers live, what they eat, and how they look, we need to first locate beavers and observe them. In this case, we depend upon certain characteristic features of beavers in order to identify them as such. As a consequence, then, the fact that they have these properties cannot be something that we discover. For, when we identify an object by means of one of its properties, we cannot possibly *find out* that it has that property – or even that it does not have this property. We can discover that beavers fell trees, but not that beavers are beavers.

The most basic way of identifying physical objects is by virtue of their position in space and time, for instance as *this item here*, or *the table that*

was here before. Again, everything that answers to the description ‘the table that was here before’ will necessarily be the table that was here before. If we refer to an object by means of its position in space and time, we cannot possibly find out that it has or does not have this position. In this sense, we know *a priori*, before looking, that the thing in question, if it exists, occupies this position.

This peculiar feature of empirical objects, that they are necessarily located at some certain position in space and time, is not something that we can learn from experience. Rather, according to Kant, we know this before we ever experience any such object since we must know it in order to experience any empirical object whatsoever. Kant writes that it is ‘the matter of all phenomena that is given to us *a posteriori*; the form must lie ready *a priori* for them in the mind, and consequently can be regarded separately from all sensation’ (Kant, 1781, B34). The forms of intuition are space and time, and since we do not learn by sense experience that empirical objects occupy spatiotemporal positions, there may be an entirely formal discipline that is concerned only with space and time.

The formal content of an experience of an empirical thing, then, is its *a priori* content in the sense specified above; it arises from the forms by virtue of which we identify an object before being able to investigate and describe it. When we refer to something as an empirical object and claim that it has a certain color or weight, we know *a priori* that we are talking about a thing in space and time, and claim to know *a posteriori* that it has this specific color or weight. That the object is located somewhere in space and time follows from the way in which we must necessarily refer to it and, thus, belongs to the formal content of our experience. The material content of our experience of an empirical thing is the information that we gather by experience: that it has this specific color or weight.

So far, the distinction between the formal and the material content of an experience may appear to be entirely relative to the way in which we come to identify a given thing. We may identify something as a rodent and find out that it is also a beaver, or we may identify something as a beaver and find out that it is a rodent. Likewise, it seems that we may refer to something *a priori* as an item that is located at the North Pole and find out that it is white, or refer to it *a priori* as a white item and find out that it is located at the North Pole. It seems to depend entirely on us which of the bits of knowledge are *a priori*, that is, what characteristic we use in order to single out the object, and which bits of information we then gather *a posteriori*, on the basis of observation.

If this is true, it would seem that we may turn the formal content of any experience into the material content of another experience and vice versa. But although this is possible for *some* kinds of experience, such as the beaver/rodent one, it is not always possible, since there are features that we must presuppose in order to identify any object. For instance, since the most general and basic way of identifying physical objects is by means of their spatiotemporal position, space and time are forms of objects about which we may have *a priori* knowledge at the most general level. We cannot really refer to a thing as a white item without knowing, at least, where it is or was located at some time. We may ask where the white item that was in Alaska is now, and answer that it is now to be found at the North Pole. In any case, we have already identified the object by means of *one* of its spatiotemporal positions; thus, we need to understand space and time in general before being able to identify any spatiotemporal object.

Yet, although in some sense, we do experience that physical objects are in space and time, this is not something that we could ever find out about them *through* experience. In order to find out anything about a physical object, we first need to locate it somewhere in space and time. Thus we never find out by observation that a thing is in space and time. In this sense, all our experience is shaped by the forms of space and time, and space and time are *introduced* by us rather than *given* to us.

Besides the forms of intuition (space and time), Kant claims that there are also *a priori* forms that our understanding introduces. For instance, whenever we unite two representations in a judgment, we must unite them in one of three ways: either one of them is a feature or attribute of the other; or one of them is a cause of the other; or both are independently and simultaneously existing entities. In any case, we apply a concept *a priori* to the representations that we combine in order to identify what they represent as real objects in the world.

Note that when we unite two representations – for instance, as cause and effect – we may be mistaken. That we apply the concept of cause and effect before being able to refer to a real object does not mean that there necessarily is such an object to which we refer. It may well be that we unite two representations under this concept in order to refer to an object, but that there is no such object. In such a case we will have applied a concept *a priori*, but in vain.

That all our experience is shaped by certain general forms which all possible objects of experience must have does *not* mean that we *construct* reality; this is a popular misconception about Kant's philosophy. We do not

bring it about that objects are in space and time when we locate and identify them as being in space and time. We use the forms of intuition and understanding in order to capture what is there, in such a way that whatever gets captured will necessarily have certain properties; namely, the properties by virtue of which we captured it. But we did not cause it to have these properties, and there might have been nothing that has these properties. In this sense, space and time and the *a priori* concepts are the forms by means of which we acquire experience.

5. *Kantian Formal Ontology*

According to the Kantian conception of the contrast between matter and form, formal ontology should be taken to be concerned with the pure forms of intuition and understanding; that is, with the way in which we must determine any object *a priori* before investigating or observing it. Its subject matter, then, will not be concrete objects, but the forms by virtue of which any experience may relate to an object. These forms will be the forms that all things have insofar as they exist. Kant claims that we can study these forms by investigating the ways in which we identify objects.

When extracting such a notion of formal ontology from Kant's writings, some qualifications are in order. Kant does not use the epithet 'ontology' to designate the study of the most general features by means of which we identify objects. Rather, he dismisses traditional ontology, identifying it with a futile attempt to say something about things that no finite rational being could possibly experience. Even to say that there may be such things, and to call them 'things,' is too much. He suggests that we should focus, instead, on our experience and on objects insofar as finite rational beings are able to experience them.

Modern ontologists, who certainly do not want to talk about objects that no finite being could possibly experience, may react to Kant's dismissal in two ways. One way is to argue, against Kant, that we do not need to reflect on our cognitive capacities in order to identify the basic structures of the world. It is true that Kant's emphasis on our cognitive judgments, and his claim that we can investigate the basic structure of the world by reflecting on our cognitive capacities, has led to the popular misunderstanding that Kant holds the world to be merely a construction of concepts. This (patently bizarre) thesis is often mistakenly labelled as 'Kantian,' both by its adherents and opponents. This misunderstanding of Kant is by no means benign, and has yielded some potentially disastrous results in modern

fields, such as information science, which seek to apply methods of formal ontology to improve the way in which information is collected, stored, and disseminated.

The other possible response to Kant's philosophy is the one adopted by Husserl, the founder of formal ontology. Roughly, Husserl continues the Kantian enterprise of investigating the basic structures of the world as it is experienced by us, but abandons his idea that there might be a reason for even speaking about anything other than the phenomenal world. Kant seems to suppose that there is a way in which the things that we experience are *in themselves*, that is, apart from all our possible experience. At the same time, however, he claims that we should not ask how things are in themselves. Husserl continues this line of thought. Since we cannot know anything about things that can in no way be experienced, we do not know whether there are such things. We do not know whether they are indeed *things*, but we also do not know whether the opposite of any of this is true. We should not try to say anything about what we cannot possibly know. But it also makes no sense to say that there are things that we do not know, or that we have no access to reality. To say that something is a thing and that it is real is already a positive claim. We cannot use the concepts 'thing', 'object', 'reality', and 'world' for the radically unknown. Rather, the world is precisely what we experience, and in this world there are real things and objects.

Insofar as ontology studies the most general features of *this* world, it is not subject to Kant's critique. That is, when Kant reflects on the phenomenal world that we experience, he already deals with the only real world that there is. For this reason, he believes that the results of our reflection on perception and experience are able to tell us what the real world is like. When seen in this light, his so-called transcendental philosophy is, in fact, the same as formal ontology. The most general structures of the world as we experience it are also the structures of the world as it really is.

This second strategy of dealing with Kant is admittedly revisionist in that it turns Kant into an ontologist, in spite of his own dismissal of traditional ontology. But the difference is, in fact, only verbal and can be traced back to two different uses of the word 'world'. According to Kant, traditional ontologists study the features of an allegedly external world that is inherently unknowable. Since we cannot, by assumption, know anything about such a world, Kant argues, it is already too much to assume that it deserves the title 'world'. For Kant, the real world is the realm of objects

that we may directly experience, and it is shaped by the general forms of space, time, causality, and so forth. Hence, when Kant declared his antipathy to ontology, he was referring to a discipline that advanced judgments about a world which is inherently unknowable, not the world as he understood it, consisting in the phenomena experienced by cognitive agents. For Husserl, ontology is concerned with the real world in Kant's sense. Kant was by no means an ontologist as he understood the term, but he was as Husserl did, and as we do.

Formal ontology studies the most general features of real objects by reflecting on the forms by virtue of which we identify them. Now, it is important to note that although formal ontology in the specified sense proceeds by reflecting on our experience, it is *not* a theory of our experience. In particular, formal ontology is not about concepts. By being formal according to the Kantian sense of 'form', ontology does not turn into a kind of psychology or epistemology, and it is certainly not the study of how a particular language or science conceptualizes a given domain. Any inquiry of this sort would have to rely on empirical knowledge about particular states of affairs; but we have already pointed out that ontology is not another special science. By the same token, it is not the study of such a special object as our experience of the world. Formal ontology is not directly concerned with particular objects of experience, nor does it have experience as its object. It is concerned with the forms that all possible objects of experience must have, and it proceeds by reflecting on experience. But to reflect on experience is not simply to form a judgment that has an experience as its object. Rather, it is to form a judgment that reflects on the way in which the experience relates, in turn, to *its* object. Thus, the object of a reflection on experience is more precisely the *relation* of experience to its object. Formal ontology is the study of how we must relate to objects before being able to investigate or describe them. It is about what it means for an experience to have a real object and what it means for a thing to be the object of an experience (see Stekeler-Weithofer, 2000, 78).

We may distinguish three kinds of inquiry. First, there are special sciences, such as biology, which study empirical things and employ concepts like 'thing', 'beaver', or 'cause' in order to describe them. Second, there are psychological or epistemological sciences, which describe and study concepts themselves by means of other concepts of a different nature such as 'concept', 'desire', or 'truth'. Finally, there is formal ontology, which studies the relation between concepts of either kind

and that to which these concepts refer. This relation is itself neither a thing nor a concept; hence, ontology does not study things or concepts directly; it studies them only indirectly, by addressing the relation that holds between them and that to which they apply. As a slogan, formal ontology does not study objects, but only the objectivity of objects.

6. *Conclusions*

The results of the preceding section lead us right back to the medieval doctrine of the transcendental attributes. We have said that formal ontology does not study objective features (or features of objects in the world) directly, but only their objectivity in and of itself. If this contrast is a genuine one, objectivity cannot be another feature of an object. And, in fact, this is what the doctrine of transcendental attributes claims. Being an object is, first, something that applies to everything that is. Everything that exists is also objective, that is, it is a possible *object of experience*. Second, *being an object* is a transcendental attribute in the neo-Platonic sense, as used by Pseudo-Dionysius; objectivity surpasses not only the boundaries between the categories, but also the boundaries of reality itself. *Being an object* is not an objective feature that could be the direct object of an experience. This means that we can study the objectivity of objects not by studying objects and their features, but only by reflecting on the relation of our experience to its objects.

We can now return to the question concerning the sense in which it is possible for formal ontology to avoid employing demonstrative expressions. It should be clear that formal ontology must include a formal treatment of space and time. But space and time, it was argued, cannot be studied without employing demonstrative expressions. However, such expressions seem to introduce material – that is, particular features – into discourse, whence it seems that there could be no purely formal ontology of space and time. What we can see more clearly, now, is the sense in which formal ontology need not depend on the *use* of demonstrative expressions in order to refer to any particular thing. Formal ontology proceeds by *reflecting on* the use of such expressions, *without using them*. It may study what it means to be or have a particular object; but, in fact, it does not refer to any such object. To reflect on an experience is to advance a judgment that relates the content of the experience to its source. When we reflect on a judgment of experience that uses a demonstrative expression, we need to understand how such an expression works, but we

need not repeat its use. A formal ontological judgment, thus, may indirectly presuppose an understanding of how demonstratives work, but it need not depend on particular or empirical facts about the things to which such expressions refer.

We can now also see what Husserl means when he describes formal ontology as the ‘eidetic science of the object as such’. Formal ontology does not make particular, empirical, claims about concrete things. In this sense, it is a formal (= *eidetic*) discipline. Further, formal ontology is a reflective discipline about the form of objectivity, that is, about what it means for something to be the object of a possible experience (Husserl, *Formale und Transzendente Logik* §38, *Husserliana* 17, 112). When Husserl says that formal ontology is the science of the object *as such*, he does not mean that it is about the object in itself *as it really is*, since *every* science should be about its object as it really is. Even a science of fake objects would be about these objects as they really are. Fake guns, for instance, really exist, and a science of fake guns should study them insofar as they really exist. So, formal ontology is not special in that it studies objects insofar as they really are. Further, by speaking of the object as such, Husserl does not mean either that ontology studies the object *apart from our knowledge of it*. We cannot study anything apart from our knowledge of it, because studying something is the process of getting to know it. Husserl, instead, uses the phrase ‘as such’ in its most straightforward and original meaning. ‘X as such’ simply means ‘*X insofar as it is X*’. Formal ontology studies objects insofar as they are objects. ‘*Object*’, however, is a relative term, as something is the object of something else. Formal ontology is about objects of possible experience insofar as they are objects of possible experience.

Chapter 3: A Primer on Knowledge Representation and Ontological Engineering

Pierre Grenon

1. *Introduction*

Suppose you want to understand the world, or a portion of the world – for instance, how a post office works – because you want to do something in the world (say, you want to send a letter to your grandmother) or because you want to know how the world will be if certain conditions obtain (say, you worry whether you will ever be able to send another letter to your grandmother should your neighborhood post office close), or just for the sake of understanding (simply because you are interested in post offices). The field of knowledge representation aims to provide computer support for doing precisely this sort of thing, namely, understanding the world or reality. At first glance, there appears to be no reason to expect that working with computers is any different from working without them. Computers are tools. Sometimes they perform well, but it is unlikely they will perform better than you prepare them to; they often perform worse. The point, however, is that there can be reasons for poor performance which are not due to limitations of computers. Knowledge representation can make computers better tools by being serious about representing the world.

If you take this endeavor seriously – that of understanding the world – you need some basic commitments which roughly amount to acknowledging (1) the reality of the world and (2) your own reality as part of the world. Let us call these the *basic principles of realism*. Also, it would be helpful if you have (3) a positive attitude toward your capacity to understand the world, and that you take the view that you are able to know something about reality, if even roughly or approximately. Let us call this the *realist principle of knowledge*. Combined, these three principles amount to the assertion that there is a reality which may be known (even imperfectly), and that knowing is a relation between a knower (you), and reality or a part thereof (the object of knowledge, such as a post office).

It is striking, then, that in knowledge representation there is a widespread attitude which tends to contradict this basic methodology. The present chapter is intended as a prophylactic against this attitude. Our

intention is to warn against denying, or ignoring, realist principles by highlighting some of the problems to which such attitudes may lead. As we will see, whereas the commonplace motto in knowledge representation is to represent convenient conceptualizations (i.e., representations) of the world, the basic motto for a sound knowledge representation must be that it is the world itself that ought to be represented. Perhaps this comes across as a platitude, and all the better if it does. But it is something that is too easily and too often dismissed as such, even though it should be constantly borne in mind.

I will begin with some non-technical preliminary considerations about knowledge representation and examine the role of philosophy in knowledge representation against this background along two dimensions. The first dimension concerns how philosophy itself is a tool for representing knowledge, whereas the second concerns the philosophical foundations of a methodology for knowledge representation. While presenting this second dimension, we put forward a methodological approach which adheres to the principles of realism mentioned above. I use these principles as representative of a philosophical position opposing other positions which could be called 'conceptualist'. In the same way that the methodology I propose is tied to the realist philosophical positions (that is, somewhat loosely and maybe only intuitively), the methodological approach I call 'traditional' is tied to conceptualist positions. To some extent, the views I name when discussing the traditional approach are mere reconstructions of positions, philosophical or methodological, which are only implicit in practice. These views are not established or structured doctrines. The reconstruction is needed in order to make a number of problems easier to grasp, problems arising from certain practices and outlooks which ignore or deny realist principles. Within these practices and outlooks, I select three specific methodological suggestions that embody in their own ways dreadful positions. Toward the end of the chapter, I formulate a more positive account of the methodological outlook needed for knowledge representation. Rather than aiming at firm guidelines and detailed recipes, this presentation of a realist methodology for knowledge representation aims to offer a taste of the mindset required for using and developing knowledge representation technologies in a sound way.

2. Preliminaries: Understanding and Representing the World

2.1. Non-technical Knowledge Representation

There is, then, the world, you, and your desire to understand the world. Fulfilling this desire is likely to involve some of the following steps, which might have to be taken recursively:

- (1) looking at the world (e.g., going to the post office and observing);
- (2) gathering facts about the world (e.g., observing some people enter with papers in their hand and leave without them, some other people behind a counter taking these papers from the former people, marking the papers and receiving money);
- (3) representing the world (e.g., taking notes of all of the above, finding some way of making perspicuous statements about what you have observed);
- (4) conjecturing the presence of sophisticated structures and the existence of other entities which may account for the facts at hand (for example: apparently the population in a post office is divided into two groups, those who pay for handing papers and those who handle the papers. Since what the latter do to the papers must be significant, perhaps there is something which makes a difference to the papers before and after the interaction, such as the stamps);
- (5) validating conjectures through inferences and experimentation (e.g., preparing a paper to send to yourself, going to the post office, verifying that there is a fee and how much it is, sending the paper, waiting and comparing the paper you receive to the paper you posted);
- (6) inferring further structures and the existence of additional entities based on valid inference patterns and the facts at hand (for example, apparently, some people deliver mail; since those you gave your mail to spend their day at a counter, there must be other people doing the delivery, and there must be a complex organization behind this).

At each step, something could go wrong which might have repercussions on the understanding you gain about the world. You might arrive at the post office while it is closed. If you do not know about hours of operation, this might lead you to infer that post offices are often or always closed. You might be overwhelmed by details or be too sensitive to the specifics of the post office you visit. If you only visit a post office in Saarbrücken, Germany, for instance, you might infer that only German is spoken in all post offices. Conversely, you might over-generalize and infer that any place with a counter is a post office. If your facts are not right or if your inferences are shaky, you might end up with quite an odd vision of the world.

One way to avoid such problems is to make sure that you get your facts straight and that you make proper inferences. This can take a great deal of care and effort, as in any science. If you are not a scientist or if you do not have the resources to spend on scientific research, you will want to ask somebody who might have the required knowledge, or to read their books. More competent people will be more helpful. The more a person knows about a domain or the more reliable her knowledge of the domain – given some criterion of reliability – the more competent she is in this domain. In increasing order of competence regarding post offices, you can ask a passerby, a post-office customer, a post-office employee, or an expert commissioned by the post-office company to design post-office regulations. Similarly, in increasing order of competence regarding the geography of Germany, you can ask a passerby on the street of another country, a passerby on a street of Berlin, a German geographer, or a geographer specialized in the geography of Germany.

These sorts of resources can improve the breadth of your factual knowledge. As your knowledge increases, it will become critical to have a way of storing it accurately and accessing it easily. It is not convenient to rely on people, because they are generally not readily available to answer your queries. Relying on paper documents might become an issue as your sources increase in size and number. You will also have difficulties finding precisely what you need in, possibly, massive amounts of irrelevant material. These are very simple and practical motivations for using technological help in storing, retrieving, and sharing knowledge.

Our concern is with several aspects of factual knowledge, namely, its quality, its efficient management, and what one can do with it. In addition to informants, one might want to turn to people with good reasoning or inferencing capabilities who can analyze one's data and extract new

knowledge from it. You might have limited knowledge and not be too sure of what else this knowledge allows you to take for granted. To a large extent, inferencing capabilities and competence, or expertise, in a domain are *prima facie* independent of one another. The sorts of inferencing relevant for us are not always very sophisticated, as we will discuss in an example below. However, they are technical and can be laborious; even in simple cases, it is a little like following an overly detailed recipe. The theory of these tools is provided by logic. In many cases, machines are very efficient at performing certain of these reasoning tasks. Here again, we find simple and practical motivations for using technological help, this time in order to manipulate, transform, and analyze the data or knowledge at hand.

2.2. Machine-Based and Formal Knowledge Representation

Looking to machine-based assistance to achieve this sort of goal brings one to the field of *knowledge representation*, which stands at the junction of the larger disciplines of *artificial intelligence* and *knowledge management*. Artificial intelligence, itself, is a field of computer science whose purpose is to get machines to perform tasks usually performed by human beings such as, in the present context, making inferences. For its part, knowledge management aims to make knowledge accessible, manipulable, and sharable, and may be seen as an attempt to produce efficient and re-usable tools for the understanding and manipulation of human and machine-processable knowledge.

A typical software solution for using machines to perform these sorts of tasks is an *expert system*. Expert systems consist of software dedicated to performing the tasks that a human expert would perform. An expert system contains three parts, namely, a *database*, an *inference engine*, and a *user interface*. Here, we are interested in the first two; the user interface is irrelevant, for our purposes. Databases are used to record facts about the world. An example of a medical fact is that there are streptococci in your throat, while another example is that the strand of streptococci in your throat is one whose presence in your throat is pathological. An example of a geographical fact is that Saarbrücken is west of Leipzig. Another example – maybe only a borderline geographical example – is the fact that family Gnomsfreunde in Saarbrücken harbors an impressive collection of garden gnomes in their yard. It is facts such as these that we record in databases. In the present context, databases which record such facts are

called *knowledge bases*. Nowadays, knowledge bases are not only parts of expert systems dedicated to a particular domain, there are also large multipurpose knowledge systems dedicated to multiple domains, or even developed to act as potential universal expert systems; for instance, the Cyc system (Lenat and Guha, 1989).

The inference engine is the part of the software that enables *inferencing*. Inferencing is the process of eliciting facts not recorded in the knowledge base, on the basis of two sources. These sources are (a) facts which are recorded in the knowledge base and (b) *rules of inference*. An example of a rule is, ‘if a family collects garden gnomes, this family decorates its house for Christmas’.

Suppose that we have a knowledge base which contains one fact and a rule:

Fact: The Gnomsfreunde family collects garden gnomes.

Rule: If a family collects garden gnomes, this family decorates its house for Christmas.

Running the expert system, we could infer from the fact and the rule the fact that:

The Gnomsfreunde family decorates its house for Christmas.

Using rules with an inference engine is a way of transforming the knowledge base so that it presents a finer and more explicit picture of the world than that provided by the raw set of facts which it initially contains. In a way, rules themselves contain knowledge; thus, the sum total of the knowledge in an expert system is the union of the knowledge base and the rules used by the inference engine.⁸

From now on, I will take the standpoint of the builder and maintainer of a knowledge base, who is often referred to as a *knowledge engineer*. The core task of the knowledge engineer is to put knowledge into computer-

⁸ The rules we are dealing with are factual (it is a fact about reality that these rules obtain, and they describe reality or relevant portions thereof). There are also logical rules of transformation which have nothing to do either with the domain or the factual knowledge at hand. For instance, the deduction above uses the logical rule of *modus ponens*: *A* (premise); If *A*, then *B* (premise); Therefore, *B* (conclusion). Going more into the details would be tedious for the unfamiliar reader and mostly irrelevant for the familiar one.

processable form. Typically, this is done by imposing a formalism upon the data, thus enabling it to be stored and manipulated. At this level, we do not need to distinguish between facts and rules, for both are of interest to the knowledge engineer. Storing and handling information are tasks that belong to the data management and information retrieval part of technological knowledge management. The resultant formal representation is often used to provide support for applications, such as reasoning (inferencing) or natural language processing, which are, in turn, sometimes considered to fall within the scope of the knowledge engineer's activity. Here, I will focus only on representational issues and not address those other activities.⁹

There are a variety of techniques and formalisms that the knowledge engineer may use (for example, compare the different ways you and your foreign language-speaking neighbor describe or refer to the same fact). I will take, as paradigmatic,¹⁰ the representation of knowledge by means of a logical formalism such as that of first-order predicate calculus. This mode of representation has the advantage of allowing, in principle, the explicit representation of the objects in the relevant domain of discourse, and a straightforward formulation of rules as statements of logical consequence.

3. *Ontology and Knowledge Representation*

3.1. *Engineering Knowledge*

A knowledge engineer deals with bodies of knowledge which include factual data and the sort of data that are contained in rules. Very quickly, it becomes obvious that these bodies of knowledge need to be given a structure, not least for reasons of efficiency and reusability. Because the knowledge engineer is engaged in manipulating and structuring knowledge, her activity is shaped *inter alia* by philosophical assumptions which underlie her adopted methodology – a truism where any activity

⁹ Actually, natural language poses a problem for knowledge representation. But, aside from the sheer difficulty of natural language processing (parsing of text via knowledge representation or generating text on the basis of formalized knowledge), it poses a problem rather indirectly and methodologically, through the more or less deliberate reliance on natural language phenomena when devising a representation of knowledge. We will come back to this when discussing what I shall later call *linguisticism*.

¹⁰ See Bibel, *et al.*, 1990, for a somewhat dated, but remarkable, introduction to the field of knowledge representation and its techniques.

such as representing the world is concerned. But her activity is also directed by presuppositions inherent in the specific domain, or implicit in the structure of the framework in which the formalization of information is conducted. To see this more clearly, let us separate out the various tasks that the knowledge engineer performs.

The work of representation begins after a body of knowledge has been acquired from the preliminary stage called *knowledge acquisition* or *information gathering*. Sometimes it is the knowledge engineer herself who gathers the facts, other times she receives her data from an external source. From this point onward, the activity can be broken down into three somewhat overlapping main tasks:

- (1) feeding the knowledge base,
- (2) improving the existing framework,
- (3) formalizing the knowledge of an expert in the pertinent subject matter.

Feeding the knowledge base means recording facts, e.g., that Saarbrücken is a city in Germany or that patient Lambda consulted Doctor Mu on a particular date. Unless this is done from scratch (i.e., the knowledge base is empty), there will already be a way to represent facts which is, more or less, specific to the system used. This means, roughly, that the language for knowledge representation might already be developed to some extent and will therefore impose constraints on which facts it is possible to represent and how to represent them. In such cases, the work of the knowledge engineer is shaped by the structure of the pre-existing framework. Take, for instance, the fact that Saarbrücken is a city. If the framework countenances a type of entity to which cities belong, representing this fact might be a matter of predicating 'being a city' of Saarbrücken, yielding 'Saarbrücken is a city'. Alternatively, if the framework associates entities such as cities with something we may call a geopolitical status, we could express a relation between Saarbrücken and the object named city, yielding 'Saarbrücken *has geopolitical status* city'.

It might be the case that the framework in which the knowledge engineer is working is not completely suitable for adding some of the new facts. This could be due to a lack of vocabulary; for instance, perhaps there is no way of expressing that Mu is a doctor. It could also be because the vocabulary that the knowledge engineer used to gather facts is already used by the framework in a conflicting way. This would be the case, for

instance, if everything in the framework which deals with medical patients were geared toward the veterinary domain and that, in particular, only a limited range of non-human animals could show as patients. At this stage, there may be a need for improvement or refinement of the existing framework itself, or of the way the knowledge engineer records the facts for herself. Issues regarding lack of vocabulary are usually rather benign. Generally, they can be resolved by extending the vocabulary used in the system. So, for instance, if you have no term for speaking of a patient in the language used by the system, or if the closest term does not precisely fit the intended use when representing a particular fact, you just add the missing term. You might have to do more work in order to amend and polish the existing structure (for example, removing the inadequate vocabulary or correcting a possible ambiguity by adding a suitable generalization). More problematic, however, are issues which have to do with limitations that are intrinsic to the language of representation. This sort of problem is one which suggests that the initial language is not suitable for representing the kinds of facts in question, in which case it might be necessary simply to opt for another language.¹¹ The only other means of overcoming such limits on expressibility is to rework the facts themselves. However, such activity is less than innocent when the goal is representation of the facts.

These first two tasks constitute, in practice, the extent of the knowledge engineer's activity. The third task, formalizing the knowledge of an expert in the pertinent subject matter, can be conceived as a way of feeding the knowledge base. But this requires that the system be able to allow for the

¹¹ Description logics (DLs) are a family of knowledge representation languages – fragments of first-order logic with nice computational properties – which can be used to represent the terminological knowledge of an application domain in a structured and formally well-understood way. Where given sorts of statements are not readily expressible in a DL language, this does not mean that the language is powerless but, rather, that one will need to find some *ad hoc* way of representing the knowledge that needs to be conveyed. This can lead in turn to non-trivial alterations of the ontological resources you had in view. Trying to preserve ontological resources and the integrity of an ontology while using languages with expressivity constraints is one important research activity in ontological engineering (see, for example, Grenon, 2006). For a less technical example, consider the different degree of user-friendliness of a language containing only adjectives, nouns, and a copula for attributing the former to the latter, with that of a language containing also verbs, adverbs, and prepositions. Or, consider again the relative merits of a communication system based on drawings against those of a full-fledged language.

formal representation of the expert's knowledge, so it might involve the second task as well. The knowledge contained in a knowledge base can be made of simple facts (e.g. Saarbrücken is a city in Germany) or rule-like knowledge, such as that all the people who reside in Saarbrücken reside in Germany. For instance, the ability to record the fact that somebody resides in a geographical region is required in order to record a rule such as that all people who reside in Saarbrücken reside in Germany. This is why new vocabulary sometimes has to be introduced before knowledge can be represented accurately. It is also worth noting that aligning new vocabulary with the existing one is a way of adding more facts and rules (that is, knowledge) to the knowledge base; for instance, expressing that residing in a region involves having a dwelling in that region.

What, then, is the knowledge engineer doing when she is building a formal vocabulary suitable for the representation of facts of a certain sort? She may be doing two things: (a) trying to fit the facts to a representational schema, or (b) trying to tailor the representational schema to her intended representation of facts. Doing the latter requires that she process and analyze the facts. It requires that she try to identify the structure of each fact, as well as what entities the fact involves. This is, roughly speaking, building a theory. Theory-building along these lines typically proceeds by *generalization* (e.g., all post offices have customers and clerks) and *abstraction* (e.g., there is a kind of entity under which clerks fall, and another under which customers fall, and each of these kinds is associated with properties reflecting powers and abilities of the relevant people). Sometimes building these sorts of theories about the world, or a domain of reality of interest to the knowledge engineers, requires making use of metaphysical and ontological insights. We will see what this means now.

3.2. *Philosophical Ontology*

Philosophical ontology is a branch of philosophy concerned with the question of what there is. More specifically, it focuses on determining what entities exist in the world and what the categories they fall under. (See Chapter 2 of this volume.) The product of an ontological investigation is typically built around one or more backbone taxonomies or hierarchies of categories. A typical ontology would include, for example, taxonomies

with the categories of *substance* and *property* as topmost nodes.¹² In addition, an ontological inquiry will provide an account of the relations between entities, and of the structure of the world at a high level of generality. For instance, the ontology will represent the relation (sometimes called exemplification or inherence depending on the type of ontology at hand) between substances and properties. Philosophers then discuss whether, given a proposed category, there are actually entities which fall under it and devise different assays of the subcategories needed. Some may deny altogether the existence of a category – for example, that of substance – claiming that the entities alleged to fall under this category are, in fact, entities of another kind (for instance, instead of being substances with properties, they are merely bunches of properties, sometimes called bundles).

The picture becomes slightly more refined, however, when we consider a distinction between at least two kinds of ontological inquiry made popular by Husserl (see, for instance, his 1931). On the one hand there is *formal ontology*, which conducts analysis and produces theories of a domain-neutral sort, theories of *forms* (for example, of part-whole relations, number, and so on). On the other hand there is *material* or *regional ontology*, which is the ontology of some specific domain or material region (for example, of mind, behavior, society, and so on).

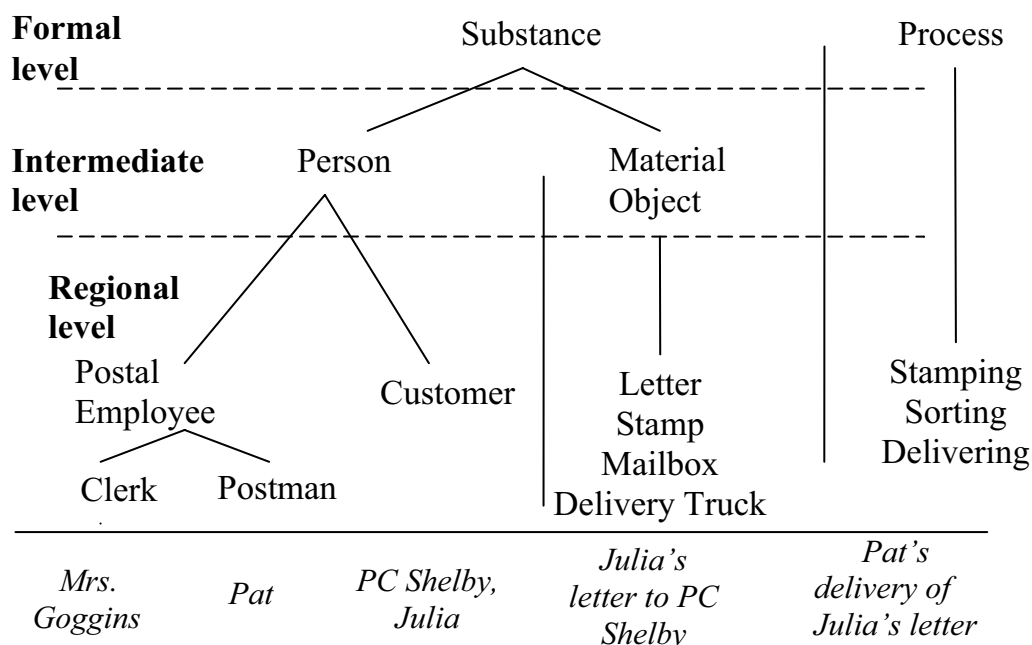
Consider the example of an ontological inquiry in the domain of post offices. This is a domain-specific inquiry, in that we are looking for the ontology of what is going on in a (typical) post office (rather than of what sorts of post offices there are). You might find that a post office has: a clerk, Mrs. Goggins; a postman, call him Pat; and a number of customers, such as Julia Pottage or PC Arthur Selby. One way of answering the ontological question (‘what is there?’) is to say that there are entities which fall under the following kinds, types, or categories: the category of clerk (Mrs. Goggins), the category of postal delivery agent (postman Pat), and the category of customers (Julia Pottage or PC Arthur Selby). This is not a very sophisticated answer – in particular, because these people are not *just* clerks, mail carriers, or post office customers – but it will do for the purposes of this simple illustration. In addition, there are a number of kinds of *activities* (things these people do) and relations between these people: Mrs. Goggins sells stamps and sorts the mail; Pat delivers the mail; PC

¹² ‘Substance’ is a term for the category under which those entities fall that may be characterized. ‘Property’ is the term for the category of entities which characterize substances.

Selby and Julia buy stamps, post mail and, perhaps, receive mail as well. So you may add the categories of *stamp selling*, *mail sorting*, *mail delivery*, *mail posting*, and so on to your ontology. Of course, more is needed; for instance, a category of *stamp* and a category of *mail object* with subcategories of *letter* and *package*. Examples of relations between the entities are *delivering mail to*, *selling stamps to*, *buying stamps from*, and so on.

All of these categories and relations are domain-specific; if you extend your ontology by generalization (a clerk is also a person) or by comparison (there are other domains in which some person buys some object from another person), you progressively arrive at less and less domain-specific considerations. You finally reach the level of domain unspecific (or, domain-neutral) considerations in which you have, say, categories of substance, property or quality, process or event, and relation. This constitutes roughly what Husserl calls the level of forms. The categories at this level can allegedly be applied and specialized in more restricted domains. In between the most specific level and the most general level you consider, there can be any number of intermediate levels. Intermediate levels are pertinent to more than one domain-specific level, but not to all.

Figure 1: *Ontological Levels*



This picture is presented in Figure 1: the dashed lines separate categories into levels; the continuous horizontal line separates categories from

examples of entities that fall under them, also called instances of those categories; and the lines between categories stand for subsumption (i.e., the category below is more specific than that above in the sense that all instances of the lower one are instances of the higher one, but not vice versa).

3.3. *Information Science Ontology*

When the knowledge engineer elaborates a theory of a domain and designs a system of categories together with the properties and relations which characterize the entities belonging to that domain, she is building an ontology. Building an ontology is, in the first place, a technical activity and not necessarily one that involves philosophical craftsmanship. But the threshold is very easily crossed. The knowledge engineer becomes an ontological engineer¹³ as soon as she performs a philosophical analysis of the content and a shaping of the infrastructure of the knowledge representation system, in the light of metaphysical/ontological theories, or in a way that is inspired by such theories. Such situations occur, for instance, when she wonders whether the fact that Saarbrücken is west of Leipzig involves one entity and its property (Saarbrücken, being west of Leipzig) or two entities and their relation (Saarbrücken, Leipzig, being west of). It occurs when she wonders whether these entities are continuants (entities which change over time while retaining something of their identity – as common sense would have persons and maybe cities do) or whether they are more like processes or events (entities which unfold in time such as rugby games).

In information science, however, the term ‘ontology’ is used in a multiplicity of ways; a fact which, over the past decade, has generated – and continues to generate – a flood of conflicting publications on the way the word should be used. In effect, the term ‘ontology’ applies to virtually any structure resembling, to some extent, a set of terms hierarchically organized which may be put in a machine-processable format. In increasing order of sophistication, the reference of the term ‘ontology’ in information science may include:

¹³ The distinction between knowledge engineering and ontological engineering essentially amounts to a division of labor. This distinction was made already by Russell and Norvig (1995) for whom it corresponds to the opposition between the domain-specific and the domain-neutral.

- (1) a set of terms (classes, categories, concepts, words),¹⁴
- (2) an axiomatic theory or a set of propositions,¹⁵
- (3) the content (conceived in a rather loose sense) of a knowledge base in general (or of some specific knowledge base such as that of the system Cyc).

A predominant view in the field of knowledge representation, reflected in the current use of the term ‘ontology’ in Semantic Web circles, is that ontological engineering is a form of modeling. On this view, an ontology would be a model or the description of a model.¹⁶ Indeed, some even speak of the task of building an ontology as a matter of conceptual modeling or of conceptual representation (a conceptual system would be a model of reality). This is but a specific case of a more general trend in knowledge representation which we will now discuss.

4. Trends in Knowledge Representation

In the Introduction, I said that the view put forward here is one according to which the basis for knowledge representation should be, not representations of reality, but reality itself. Let us call this view *realist representationalism*. On this view, when we are doing ontology, we are dealing with the things themselves, not with representations of them. While this principle may seem too obvious to need mentioning, the most widespread methodological stance in knowledge representation leads, in fact, to practices that oppose realist representationalism. This stance puts knowledge engineers at risk of committing the sorts of serious blunders which we address in the next section and elsewhere in this volume.

I will call the mainstream methodological approach to knowledge representation *pragmatist conceptualism*. According to this approach, the knowledge engineer’s main priority is to create a smoothly functioning knowledge base out of whatever conceptualization is provided to him by domain experts. This approach conceives the task of the knowledge

¹⁴ Such sets can be more or less structured and can be anything from taxonomies to dictionaries, vocabularies, terminological systems or even thesauri as well as semantic networks such as WordNet (<http://wordnet.princeton.edu>).

¹⁵ See the formal treatment of theories developed in Menzel’s *Ontology Theory* (2003)

¹⁶ A model, here, is probably best understood as a representation which allows a certain degree of simulation and approximation judged adequate by the modeller for the purpose at hand.

engineer as consisting *only* in that of representing others' (the domain experts') representations, so that reality falls out of the picture almost entirely.

If a situation arises in which the expert's conceptualization does not make for the most smoothly functioning knowledge base (perhaps because the logical language in question cannot express a certain kind of statement), pragmatist conceptualism licenses the knowledge engineer to adjust the expert's conceptualization to make it expressible by means of the tools he has at hand. Thus, pragmatist conceptualism frees up the knowledge engineer to bring up any kind of objects and any theoretical construct that may prove useful for her representation. The only guidance is that such conjuring logically fulfill the practical purpose of the representation. This is emblematically endorsed by Genesereth and Nilsson (1987), who sum up their position as follows: 'no attention has been paid to the question whether the objects in one's conceptualization of the world really exist.... Conceptualizations are our inventions, and their justification is based solely on their utility' (p. 13).

But what, exactly, *is* a conceptualization? This is far from clear, even among knowledge engineers themselves. Minimally, a conceptualization involves concepts and probably also their specifications. In turn, a concept is probably a thing which carries some sense or meaning. However, if any word is polysemous, it is the word 'concept' (see Chapter 4). Information science, in particular, puts far too heavy a load upon the term 'concept'. There does not seem to be any prevailing meaning, and it is rarely used with a single, coherent meaning. For example, 'concept' might be taken to be one of the following: (1) an idea or a mental representation of objects in reality; (2) a general idea under which a multiplicity of things falls (let us call these *conceptual universals*);¹⁷ (3) a Platonic¹⁸ idea existing as a perfect prototype of things in the world, but itself, in some sense, exterior to the world; (4) a class, set or collection; (5) a word; (6) the meaning of a word.

These various meanings of 'concept' are often run together in more or less subtle ways.¹⁹ The pervasive use of the term, and the running together

¹⁷ Reflected in a language by the use of a (typically monadic) predicate, they allow great flexibility in representation, but it should be noted that only some correspond to any counterparts in reality. There is no universal in the world, for example, corresponding to our general concept of a unicorn.

¹⁸ Here this term can be taken as a synonym for abstract and perfect model.

¹⁹ This is such a routine difficulty that it can even be experienced with an ISO standard such as the terminological standard 1087-1:2000, for example (see ISO).

of its various meanings, result in further confusion about what a knowledge engineer is after – what she is trying to represent – and how she should carry out her work. The distinction between things (or real entities) and the corresponding concepts is acknowledged by the knowledge engineer, if only when endorsing the claim that conceptualizations are independent of reality. In her actual work, however, things in reality are neglected to the benefit of their conceptual proxies.

It is quite clear from the passage cited above that the sort of conceptualizations Genesereth and Nilsson have in mind are made of concepts in a non-realist, mind-dependent, sense rather than entities, such as universals or kinds, existing independently of human cognition (see Chapter 8 of this volume). To be fair, the actual nature of concepts is probably of little interest to a large number of knowledge engineers. It is credible that, the closer they are to computer science than to philosophy, the more susceptible they would be to being dismissive of what they would see as mere philosophical hair-splitting. And, to some extent, this dismissal would be understandable. It is conceivable that issues for philosophers could be non-issues for computer scientists and vice versa. But the problem here is that there are practical consequences associated with whether we adopt a methodology for knowledge representation that is inspired by realism, or by conceptualism.

One problem with pragmatist conceptualism is that, in the long run, representation that is the most *useful* is actually that which is the most accurate in relation to *reality*. Now, for computer scientists, usefulness might mean that modeling is easy or that inference is fast. There can then be architectural reasons specific to given systems which, for the sake of ease in using that system and of optimizing that system's inferencing resources, lead the knowledge engineer to adopt *ad hoc* modeling solutions and simplistic misrepresentations when using that system for representing knowledge. Giving in to such practices, however, results in idiosyncrasy (see Smith, 2006). Tailoring formal representations to suit the optimal settings of a given knowledge system might be less stable in the long run (because these settings are tied to hard coded features of the system and, thus, are dependent on the level of development of the architectural components of the system) and less reusable (because, obviously, they are *ad hoc*).

It is perfectly proper that we, as knowledge engineers, should be looking for a useful account. But it is hard to imagine what greater usefulness a knowledge representation could have than to be accurate with

respect to reality. It should be clear that, if the knowledge engineer professes pragmatic conceptualism explicitly, then there is no reason to accept her representation as anything more than daydreaming or literary fiction. A suggestion to those concerned exclusively with efficiency, then, might be that, if given tools will not perform without an *ad hoc* representation, then these tools need either to be improved (so as to allow for a representation that is more adequate to reality), or discarded. In contradistinction, aiming for consistent adequacy to reality in knowledge representation is likely to result in an improvement in the stability and usefulness of the resulting framework and, all things being equal, a broader acceptance of that framework.

Another disadvantage to pragmatist conceptualism is that it yields knowledge bases which are internally consistent, but unsuited for being linked with other knowledge bases. For example, it provides us with no principle ensuring that, for instance, bottom-up (from domain-specific to domain-neutral) and top-down (from domain-neutral to domain-specific) approaches will meet in any coherent way (suppose you want to place your post-office ontology under a more general ontology of services). Nor does it ensure us that two independently built ontologies or two independent knowledge representations of the same domain will overlap or agree about even one fact because, on this position, there are no *facts* to agree about.

For a long time, too little attention was paid to a principled resolution of the problem of unifying the many different, and often mutually incompatible, idiomatic frameworks and representations developed independently by different groups or companies. This is sometimes informally referred to as the *Tower of Babel problem*, and early attempts to solve this problem tried to devise schemas to which existing knowledge representations could relate, putting in place platforms of translations. For instance, the Knowledge Interchange Format (see: logic.stanford.edu/kif/dpans.html) was first conceived of as a language to which other knowledge representation languages would be mapped, providing a central node in the net of inter-translatable languages. A sequel to this effort is the attempt under the Common Logic rubric of producing, as an ISO standard, an even more general and abstract specification of a knowledge representation language (see <http://cl.tamu.edu>). On the side of ontology, the Cyc ontology has been marketed as a potential platform for linking and comparing different ontologies. More recently, the IEEE gave its blessing to a Standard Upper Level Ontology working group (see <http://suo.ieee.org>), which aims to devise a consensus top-level ontology.

For various reasons, not all technical, there is still no ontology which is accepted as a standard in the knowledge-representation community. Instead, there are many candidates among which, for example, are OpenCyc (www.opencyc.org) and SUMO (www.ontologyportal.org), and more candidates have been springing into existence on a regular basis.

There are both pragmatic and ideological reasons explaining this state of affairs and the reluctance to work toward a common ontology. On the pragmatic side, it involves the resolution of a number of non-trivial problems (including problems of logic), and it would thus take time to come up with a decent standard. Such pragmatic considerations have sometimes been used to belittle the value of the attempt. But even if an acceptable standard ontology were successfully created, it would take more time and possibly prohibitive amounts of money to bring existing ontologies and knowledge bases up to a level of compliance. On the ideological side, what is at issue is the nearly pervasive, though often merely tacit, adherence to one or other form of conceptualism (often with constructivist or relativist leanings), which serves to render immediately suspect any candidate that is put forward as *the* standard ontology. This has the consequence that the question of the adequacy to reality of the conceptual schemas tends to be neglected. It also seems to imply that the search for an ontology of reality, rather than an ontology of its multiple representations, would bear no fruit.

The general problem of standardization has given rise to an area of research into *interoperability*, which focuses on how to manage the joint operation of distinct frameworks.²⁰ The premise is that if there is a multiplicity of conceptualizations, all should be accounted for. In effect, we end up with concurrent systems allegedly representing the same reality but, in fact, failing to do so because of presupposed *de facto* and in principle problems of interoperability. If concepts differ, what is to serve as our guide in resolving these differences? That is, if we wish to unite these competing systems together, what can serve as a *tertium quid*? The question becomes one of how to relate (or ‘fuse’ or ‘merge’) them. That such merging is a difficult problem – as illustrated, for example, by the Unified Medical Language System ([http:// umlsinfo.nlm.nih.gov/](http://umlsinfo.nlm.nih.gov/)) – reflects some of the dangers into which pragmatic conceptualism leads us.

²⁰ In the context of the Semantic Web, the recourse to namespaces, using a syntactic device marking the origin of each concept, helps as a bookkeeping device; but it eliminates only some of the confusion. It does not provide the needed ticket for interoperability and, probably, is only a short-term solution.

5. *Tendencies in Conceptualist Knowledge Representation*

There are at least three types of problematic tendencies to which pragmatist conceptualism opens the door; I will call them *linguisticism*, *algebraism*, and *subjectivism*. These tendencies echo and motivate some of the criticisms sketched above. For the sake of simplification, I present them as views in their most extreme forms. What I mean to warn against are the risks that result from making the activity of the ontologist succumb to, and crumble under, linguistic and imaginative fancies.

5.1. *Linguisticism*

The linguistic tendency rests on the conviction that natural language is already the best medium for representing reality (it is, after all, the most commonly used); thus, that the closer an ontology is to natural language, the more perfect it will be. Methodologically, but without reaching such an extreme, this translates into introducing alleged entities into an ontology on the sole basis that there are natural language expressions using terms which seem to refer to them. For instance, consider whether you agree that there truly are *sakes*, i.e. entities to which expressions such as ‘I did it for John’s sake’ refer. Consider also whether you believe that the terms ‘Santa Claus’, ‘Postman Pat’, ‘Quetzalcoatl the Aztec sky god’, and so on, designate real entities.

If natural language is allowed to dictate an ontological inquiry unimpeded, the ontologist risks making the radical claim that an ontology must stand in a one-to-one correspondence with the elements of natural language (or as close thereto as possible). A motto for this view could be ‘one word, one concept’. The main problem is that this results in an overwhelmingly rich system of concepts with no particular ontological credentials. There is a converse issue which is tied to the phenomenon of differences; for example, in lexical gaps, as between different natural languages.²¹ These lead to the proclaiming of absurdities such as that there is, for example, a French ontology and an English ontology. However, the fact that your mother tongue has no ready-made term for a given entity or kind of entity does not prevent you from using or understanding a corresponding concept or talking about the entity in question; for example,

²¹ Lexical gaps are visible when a language does not have a simple term for referring to certain entities. This can also be illustrated by the languages which lack technical vocabulary existing only in English.

by means of some more complex word formation. Above all, whether or not there is a term in a language does not determine whether or not there is something in reality to which the words or concepts correspond.

At best, natural language can serve as one clue for the ontologist, but it should certainly not be a *criterion* of the correctness of the end-result of her labors. Indeed, if everything in natural language corresponded to some aspect of reality, then there would seem to be no need for ontological structuring or conceptual modeling to begin with.

5.2. *Algebraism*

The second tendency I shall warn against can be called algebraism, which takes two forms – namely, conceptual algebraism and theoretical algebraism – depending on the objects considered: on the one hand, concepts (or general terms in a language), and on the other hand, theories. The former relates to issues of content; the latter to issues of structure.

No matter what the object considered, the common assumption behind algebraism is a view according to which a knowledge representation system ought to be maximally complete and contain all algebraically possible objects of manipulation; for example, all set-theoretically possible objects. The idea is that all combinatorial variants fall within the scope of the resultant knowledge representation system. That a total conceptualization can be viewed as a system of concepts with both an initial set and an articulating theory provides the root for the distinction between the two forms of algebraism.

On the concept variant of algebraism, each set of concepts is associated with a set of all Boolean combinations including: disjunctions, such as *dog or cat*; conjunctions, such as *red and square*; complements, such as *non-German*; and conditionals such as *if colorless then magnetic*. Further, the members are regarded as being of equal standing with the members of the initial set. Some examples of dubious objects resulting from such formal motivations are categories such as *green and number* or *horse and violin*. A related case is the postulation of a null region of space in certain mereotopologies, done to ensure that the theory behaves smoothly. Such constructs may be useful for logical purposes, but this does not guarantee them any ontological footing (see Grenon, 2003).

On the theory variant of algebraism, a theory is viewed as a set of sentences to which similar operations are applied as were applied to concepts in the above. The goal, again, is to create an end-result with a

certain kind of algebraic elegance; for example, a complete set of mutually inconsistent theories. The striving for this sort of completeness seems to have been in part responsible for the difficulties faced, for example, by the IEEE's quest for a standard upper ontology (see *IEEE* <http://suo.ieee.org>). The quest for a consensus upper merged ontology, too often, has been transformed into a process of registration of existing separate competitor ontologies.

Our dissatisfaction with algebraism arises from what we believe to be a reasonable skepticism concerning the arbitrary production of fictions for the sake of systematic neatness. We favor hunter-gatherer ontology, based on empirical evidence, rather than armchair ontology based on permutations and combinations. In practice, the latter leads to an explosion of the domain of objects (concepts or theories) which the system needs to handle, and this also leads to a problem of determining relevance and of choosing among all the theoretically possible variants. Most of all, it does not account for those natural segmentations of reality, which fall far short of algebraic neatness in domains like biology or medicine. Even if it is credible that, in some instances, different candidate theories would be retained on an equal footing, retaining an entire articulated range of combinatorially possible theories seems to be without real motivation, and to be alien to the methods of good science. Perhaps preserving the entirety of those theories which stand as alternative solutions to an as yet unsolved and critical problem might be warranted. There might be other considerations for subscribing to elements of algebraism; for instance, having to do with legacy issues (as when a system has been around long enough or used widely enough that it can be amended only with difficulty). But, even here, the combinatorial completeness demanded by the algebraist approach will not prove practically useful.

5.3. *Subjectivism and Relativism*

A third tendency could be named *subjectivism*. According to this view, the world is the product of a subject's conceptualization. In an extreme version, there are as many conceptualizations, and as many ontologies, as there are conceptualizing agents. For the subjectivist in knowledge representation, an ontology is nothing but a conceptualization, which may or may not be based on a consensus shared by a plurality of individuals. This raises the obvious problem of radical and permanent interpretation: how do we know that we understand one another, even when we speak a

common language? In other words, it takes the Tower of Babel as a premise.

This view is akin to relativism and to a position defended in philosophy known as perspectivism, according to which no human perspective (no conceptualization) has greater value than another; thus, no human deserves to be called an expert. From there, it is but a small step to claim that no human perspective is closer to the truth. And from there again, it is deceptively simple to issue a platitude such as ‘all perspectives are, if the purpose is only appropriately specified, equally useful’.

An assumption of relativism is clearly active in pragmatist conceptualism, and implicit in many approaches to knowledge representation. It is one of the underlying motivations for taking the problem of interoperability seriously. Relativism is also an obvious motivation for algebraism about theories. There is a link between linguisticism and relativism as well. Many are not afraid to go from differences in languages to differences in conceptualizations, and from there to differences in ontologies and, finally, to differences in the realities in which speakers of those languages have evolved.

6. Representation of Reality

For the reasons discussed, a line should be drawn between external reality and whatever our systems of private or collective representation might be. The focus of ontology is not the latter, but the former. But there is nonetheless some truth behind the motivations of the pragmatist conceptualist: for one thing, we do want our representations to be useful; for another, people do actually differ and disagree in their representations of the world.

How can a knowledge representation system be made rich enough to represent epistemic and cognitive facts? The language of the Cyc system is rich in that sense, including a number of suitable constructs (see (<http://www.cyc.com/cycdoc/vocab/vocab-toc.html>)). It allows expression of the fact that an agent – for instance, a doctor – will believe that a patient is in a condition of a certain type, while simultaneously allowing the diagnosis to have only a provisional nature and allowing for a separation of the true condition of the patient from the conjectured one. This sort of approach is possible in principle because, while beliefs and other cognitive stances are real, they do not inform reality or the portions of reality that

they address or pretend to address and in which they are found, but only the cognitive domain itself.

If the representation of epistemic stances requires an accurate realist account of the reality toward which these stances are directed, where shall we find objective sources of knowledge about cognition-independent reality? At this stage three maxims suggest themselves:

- i) The primary source of knowledge is reality.
- ii) The domain being represented is a part of reality.
- iii) The expert knows the reality at hand (that is why we call him an expert).

In many domains, the best experts we have at hand are scientists (natural and social). I will take the scientist as the paradigmatic expert. But, of course, in many domains we might turn, rather, to an engineer. One remark, in relation to what has been said before, is that in a scientific domain, knowledge representation would be carried out on the basis of scientific theories, but it would not be a representation of these theories. Another point, of more importance, is that reliance on science aims, among other things, to provide ontology itself with a methodology and a goal that follows scientific standards.

A realist methodology trusts experts for matters of ontology. It asserts, for the purposes that are of prime concern in the present volume, that knowledge of reality can be obtained through an inquiry of the sort conducted via application of the scientific method (that is, ‘knowledge’ in the non-philosophical sense of *very well justified beliefs*; see the Introduction for this distinction). Empirical inquiry, on this view, can furnish knowledge of the world. Now, it is a fact that our ability to engage in such inquiry evolves and progresses. It must be, therefore, that our theories and our concomitant understanding of the world can be subjected to testing and revision; it is this that secures the possibility of their progress. Indeed, there is hardly any cutting-edge theory which is not the subject of contemporary debate and revision.

This means that the ontologist who is concerned with reality must make provisions for the evolution and refinement of the views underlying her work. This is the methodological doctrine of *fallibilism*, which makes room for approximations and errors in state-of-the-art knowledge. If we are looking for accuracy, we must accept the possibility of being in error, in the hope that we may follow our inquiries toward more refined

understanding in the future. We must be ready to abandon views and introduce unforeseen elements, even if this requires us to redo laborious work.

Error is not the only source of unease for the ontologist; indeterminacy is yet another issue. Sometimes empirical inquiry remains undecided and offers incompatible, but equally credible, theories. Consider, for example, the alternative between the wave and corpuscular theories of light in the 1930s. This means that, as ontologists, our methodology must also make provisions for the possibility of maintaining equally legitimate realist perspectives on reality. This will generally occur when phenomena can be accounted for by independent and contradictory stances, not only scientific, but also philosophical ones.²²

It is important to bear in mind that this does not amount to the thesis that any view of reality is legitimate. Rather, it is a *realist perspectivalism* (as contrasted with *relativist* perspectivalisms of a more traditional sort) which asserts that, at least, some views of reality are legitimate (see Chapter 6). It is also a *realist adequatism*, which means that it denies the doctrine of reductionism, according to which the putative plurality of legitimate views of reality is not to be eliminated through the reduction of all such views to one central basic view; for example, the view of microphysics. To establish which views are legitimate, we must weigh them against each other and against their ability to survive critical tests when confronted with reality as, for example, in scientific experiments; including not only experiments in microphysics but also in biology, medicine, and other branches of science. Those concepts and conceptualizations which survive are then transparent to reality, to use the somewhat metaphorical expression of Smith (2003). More generally, we are concerned with those views that are *veridical under a given perspective* in relation to a particular domain. This means, among other things, that we can do the ontology of post offices without worrying about elementary particles.

7. Conclusion

Philosophical and ontological analysis has to be performed by the knowledge engineer in order to provide a sound basis for her knowledge representation; even when such representation is conceived as conceptual

²² Consider the tangled issue of endurance versus perdurance in the metaphysical debate over persistence in time. See Kanzian, 2008.

modeling. More precisely, knowledge representation systems ought to be conceived as representations of *reality*, and not as systems of representations of concepts, or as mere models with no foothold in the real world of what happens and is the case. For these reasons, an ontological inquiry in its philosophical sense, one which addresses reality, must be the basic methodological principle of a sound approach to knowledge representation.

A complete knowledge representation system should be able to accommodate and articulate what may be a multiplicity of legitimate views of reality. Which alternative theories or perspectives on reality are useful for the purposes of the knowledge engineer has to be established on the basis of a realist and fallibilist methodology, and this task, which comes close to the tasks of empirical inquiry, may be the most difficult and challenging to accomplish. What matters is that the knowledge engineer should bear in mind: first, that her target is reality; second, that formal simplifications, modeling tricks, and shortcuts of various other sorts, may be detrimental to the ultimate goal of accounting for reality.

Shimon Edelman's Riddle of Representation (Edelman, 1998) reads as follows:

Q: two humans, a monkey, and a robot are looking at a piece of cheese; what is common to the representational processes in their visual systems?

A: the cheese, of course. (Meaning it ain't in the head; putting it there is just as pointless as making a sandwich with a picture of a piece of brie.)

Of course, this is a metaphor; you don't take some real entity – for example, a lump of goat cheese – and put it in a database. But this is a methodological point. The knowledge base ought to be a reproduction of reality. Knowledge representation and ontology are not, strictly speaking, representations, they are *re*-presentations of reality. If the knowledge engineer were to hold fast to only one methodological proposal, then it should be this: *the world itself* should be included in a knowledge representation or ontology.

Chapter 4: New Desiderata for Biomedical Terminologies

Barry Smith

Part I: *Introducing Concepts*

I.1. *Introduction*

The terminologies used in biomedical research, clinical practice, and health information management today grew out of the medical dictionaries of an earlier era. Such dictionaries, of course, were created to be used by human beings, and the early steps towards standardization of terminologies in the 1930s were designed, above all, to enable clear understanding of terms in different languages; for example, on the part of those engaged in gathering data on an international scale. With the increasing importance of computers, however, came the recognition that standardization of terminology must go beyond the needs of humans, and it is especially in the biomedical domain – with terminologies such as SNOMED (see *SNOMED 2007*) and controlled vocabularies such as the Gene Ontology (see *Gene Ontology 2007*) – that the power of formal representation of terminological knowledge has been explored most systematically. The need for such formal, computer-processable representations becomes all the more urgent with the enormous increase in the amounts and varieties of data with which biomedical researchers are confronted, data which can no longer be surveyed without the aid of powerful informatics tools.

I.2. *The Concept Orientation*

Unfortunately, the new formalized biomedical terminologies were developed against the background of what are now coming to be recognized as a series of major and minor philosophical errors. Very roughly, the developers of terminologies made the assumption that we cannot have knowledge of the real world, but only of our thoughts. Therefore, they inferred, it is thoughts to which our terms (and our terminologies) necessarily refer – thoughts which, as we shall see, were understood as being crystallized in the form of what were called concepts.

What the term ‘concept’ might precisely mean, however, was never clearly expressed, and it takes some considerable pains to extract a

coherent reading of this term from the standard terminological literature. In fact, four loose families of readings can be distinguished, which we can refer to as the *linguistic*, the *psychological*, the *epistemological*, and the *ontological*. On the linguistic view, concepts are *general terms* whose meanings have been somehow regimented (or, as on some variants of the view, they are these meanings themselves). On the psychological view, concepts are *mental entities* analogous to ideas or beliefs. On the epistemological view, concepts are *units of knowledge*, such as your child's concept of a cat or of a square. And on the ontological view, concepts are *abstractions of kinds or of properties* (i.e., of general invariant patterns) belonging to entities in the world. As we will see in what follows, elements of all these views can be found, in various combinations, in the literature (Smith, 2004).

The most influential biomedical terminologies, including almost all of the terminologies collected together in the Metathesaurus of the Unified Medical Language System (see National Library of Medicine), have been developed in the spirit of the concept orientation (Smith, 2005a). These terminologies have proved to be of great practical importance in the development of biomedical informatics. However, the ambiguities surrounding their use of the term 'concept' engenders problems which have been neglected in the informatics literature. As will become clear in what follows, the concept orientation exacerbates many of the problems which it was intended to solve, and introduces new problems of its own.

I.2.1. The Birth of the Concept Orientation (I): Eugen Wüster and the International Organization for Standardization

The concept orientation in terminology work goes back at least as far as the 1930s, when Eugen Wüster began to develop a theory of terms and concepts which later became entrenched as the terminology standard promulgated by the International Organization for Standardization (ISO) (ISO, N.D.; Smith, 2005b). Through the powerful influence of the ISO, Wüster's standard continues to be felt today wherever standardized terminologies are needed, not least in the areas of biomedicine and biomedical informatics. However, Wüster's standard was developed for terminologies used by humans; it does not meet the requirements placed on standardized terminologies in the era of the computer. In spite of this, the quasi-legal precedent-based policies of ISO – in which newer standards are required to conform as far as possible to those already established – have

prevented adequate adaptation of standards. Even the most recent ISO standards developed in the terminology domain betray a sloppiness and lack of clarity in their formulations which falls far short of meeting contemporary requirements.

Human language is in constant flux. Focusing terminology development on the study of concepts was, for Wüster, a way of shouting ‘Stop!’ in the attempt to sidestep the tide of variances in human language use, which he saw as impediments to human communication across languages; for example (and uppermost in Wüster’s own mind) in the context of international trade. (Wüster, himself, was a businessman and manufacturer of woodworking machinery.) Since the actual human thoughts associated with language use are an unreliable foundation upon which to base any system for standardizing the use of words, Wüster’s solution was to effectively invent a new realm – the realm of concepts – in which the normal ebb and flow of human thought associated with the hitherto predominating term orientation would be somehow neutralized. Consider, for example, the way in which a term like ‘cell’ is used in different contexts to mean unit of life, a small enclosed space, a small militant group, unit in a grid or pigeonhole system, and so forth. From Wüster’s point of view, there was a different concept associated with each of these contexts. Concepts, somehow, are crystallized out of the amorphous variety of different usages among the different groups of human beings involved.

At the same time, Wüster defended a *psychological* view of these concepts – which means that he saw concepts as mental entities – sometimes writing as if, in order to apprehend concepts, we would need to gain access to the interiors of each other’s brains (Wüster, 2003):

If a speaker wishes to draw the attention of an interlocutor to a particular individual object, which is visible to both parties or which he carries with him, he only has to point to it, or, respectively, show it. If the object, however, is in another place, it is normally impossible to produce it for the purpose of showing it. In this case the only thing available is the individual concept of the object, provided that it is readily accessible in the heads of both persons.

Thus, for Wüster a concept is an *element of thought*, existing entirely in the minds of human subjects. On this view, an individual concept (such as *blood*) is a mental surrogate of an individual object (such as the blood running through your veins); a general concept (such as *rabbit* or *fruit*) is a mental surrogate of a plurality of objects (Smith, 2005b). Individual concepts stand for objects which human beings are able to apprehend

through perceptual experience. General concepts stand for similarities between these objects. Both individual and general concepts are human creations, and the hierarchy of general concepts (from, say, *Granny Smith* to *apple* to *fruit*) arises as the cumulative reflection of the choices made by humans in grouping objects together. Since these choices will vary from one community to another, standardization is needed in order to determine a common set of general concepts to which terminologies would be related; for instance, in order to remove obstacles to international trade.

The perceived similarities which serve as starting points for such groupings are reified by Wüster under the heading of what he calls ‘characteristics’, a term which, like the term ‘concept’, has been embraced by the terminology community (and, thereby, has also fallen prey to a variety of conflicting views). In some passages, Wüster himself seems happy to identify characteristics with *properties* on the side of the objects themselves. In others, however, he identifies them as further *concepts*, so that they too (incoherently) would exist in the heads of human beings (Smith, 2005b). Thus, Wüster’s thought results in an uncomfortable straddling of the realms of mind (ideas and meanings) and world (objects and their properties).

This fissure appears in Wüster’s treatment of the *extension* of a concept as well, which he sometimes conceives in the standard way as the ‘totality of all individual objects which fall under a given concept’ (Smith, 2006; Wüster, 1979). Unfortunately, Wüster also allows a second reading of ‘extension’ as meaning ‘the totality of all subordinated concepts’. So, on the one hand the extension of the concept *pneumonia* would be the totality of *cases* or *instances* of pneumonia; but, on the other hand, it would be a collection of more specific concepts (*bacterial pneumonia*, *viral pneumonia*, *mycoplasma pneumonia*, *interstitial pneumonia*, *horse pneumonia*, and so on).

Another characteristic unclarity of Wüster’s thinking is reflected in his definition of ‘object’ as ‘anything to which human thought is or can be directed’. This definition has been given normative standing through its adoption in the relevant ISO standards, which similarly define ‘object’ as ‘anything perceived or conceived’ (ISO, ‘Text for FDIS 704. Terminology work: Principles and methods’).

This ISO definition implies that ‘object’ can embrace, in Wüsterian spirit, not only the material but also the immaterial, not only the real but also the ‘purely imagined, for example, a unicorn, a philosopher’s stone, or a literary character’ (ISO, Information Technology for Learning,

Education, and Training; ISO, Vocabulary of Terminology). Given this characterization of ‘object’, we believe, ISO undercuts any view of the relation between concepts and corresponding objects in reality that might be compatible with the needs of empirical science (including the needs of contemporary evidence-based medicine). For its definition of ‘object’ would imply that the extension of the concept *pneumonia* should be allowed to include, not only your pneumonia and my pneumonia, but also, for example, cases of unicorn pneumonia or of pneumonia in Russian fiction. Of course there is nothing wrong with employing the term ‘object’ to mean, roughly, ‘anything to which human thought can be directed’. The problem is that ISO allows no *other* term which would be used to distinguish those terms which are intended to be directed towards real things and those terms which merely refer to objects in this very loose sense. Matters are made even worse by ISO’s edict that:

[i]n the course of producing a terminology, *philosophical discussions* on whether an object actually exists in reality... *are to be avoided*. Objects are assumed to exist and attention is to be focused on how one deals with objects for the purposes of communication. (ISO, ‘Text for FDIS 704’)

It is precisely such philosophical discussions which are *required* if we are to undo the sore effects of Wüster’s influence.

More recent ISO documents reveal efforts to increase clarity by embracing elements of a more properly ontological reading of the term ‘concept’, the view that concepts are abstractions of kinds which exist in the world. Unfortunately, however, in keeping with ISO’s quasi-legal view of standards as enjoying some of the attributes of *stare decisis*, this is done in such a way that remnants of the older views are still allowed to remain. Thus, in ISO 1087-1:2000, ‘concept’ is defined variously as a ‘unit of thought constituted through abstraction on the basis of properties common to a set of objects’, or ‘unit of knowledge created by a unique combination of characteristics’, where ‘characteristic’ is defined as an ‘abstraction of a property of an object or of a set of objects’. Since ‘object’ is still defined as ‘anything perceivable or conceivable’ (a unicorn still being listed by ISO as a specific example of the latter), the clarificatory effects of this move are, once again, rendered nugatory by the surrounding accumulation of inconsistencies.

As Temmerman argues, Wüster’s version of the concept orientation stands in conflict with many of the insights gained through research in cognitive science in recent years (Temmerman, 2000). His account of

concept learning and his insistence on the arbitrariness of concept-formation rest on ideas that have long since been called into question by cognitive scientists. Even very small children manifest, in surprisingly uniform ways, an ability to apprehend objects in their surroundings as instances of natural kinds in ways which go far beyond what they apprehend in perceptual experience. Thus, there is now much evidence (documented, for example, in Gelman, 1991) to the effect that our ability to cognize objects and processes in a domain like biology rests on a shared innate capacity to apprehend our surrounding world in terms of (invisible) underlying structures or powers (whose workings we may subsequently learn to comprehend; for example through inquiries in genetics).

1.2.2. The Birth of the Concept Orientation (II): James Cimino's Desiderata

By the time of James Cimino's important paper (Cimino, 1998), biomedical terminologies faced two major problems. The first problem concerned the legacy of the influential concept orientation as conceived by Wüster, which we will explore in greater depth in what follows. The upshot of this legacy was an endemic lack of precision, not only with regard to what concepts might be, but also with regard to their role in terminology work. The second problem revolved around the introduction of computers into the terminological domain. Computer-based applications rely on precision, in both syntax and semantics, in a way that human cognition does not.

In an attempt to address these problems, James Cimino introduced a set of desiderata which must be satisfied by medical terminologies if they are to support modern computer applications. In what follows, we shall argue that many of Cimino's desiderata ought to be accepted by those involved in terminology work; but only when they have been subjected to radical reinterpretation.

Cimino's principal thesis is that those involved in terminology work should focus their attentions, not on terms or words or their meanings, but rather on concepts. Unlike Wüster, Cimino comes close to embracing a linguistic rather than a psychological view of concepts. A concept, he says, is 'an embodiment of a particular meaning' (Cimino, 1998, p. 395), which means that it is something like a term that has been extricated from the flow of language so as not to change when the language does. One of his desiderata for a well-constructed medical terminology is accordingly that

of concept permanence: *the meaning of a concept, once created, is inviolate*. Three further desiderata are:

Concepts which form the nodes of the terminology must correspond to at least one meaning (non-vagueness).

Concepts must correspond to no more than one meaning (non-ambiguity).

Meanings must themselves correspond to no more than one concept (non-redundancy).

If these requirements are met, the preferred terms of a well-constructed terminology will be mapped in one-to-one fashion to corresponding meanings. (A preferred term is that term out of a set of synonyms which the terminology chooses to link directly to a definition.) On Cimino's view, a concept corresponds to a plurality of words and expressions that are synonymous with one another.

However, Cimino recognizes that synonymy is not an equivalence relation dividing up the domain of terms neatly into disjoint sets of synonyms. Often, words which are synonyms relative to some types of context are not synonyms relative to others (e.g., a bat in a cave is not the same as a bat in a baseball game). To resolve this problem, he invokes the further desideratum of *context representation*, which requires a terminology to specify, formally and explicitly, the way in which a concept is used within different types of contexts. (We will leave open the question of whether, if concepts can be used differently in different contexts, this violates the non-ambiguity desideratum.) If, however, we are right in our view that concepts, for Cimino, are themselves (or correspond in one-to-one fashion to) sets of synonyms, then concepts should thereby be relativized to contexts already. Thus, in formulating the desideratum of context representation he ought more properly to speak, not of *concepts*, but rather of *terms* themselves, as these are used in different types of contexts. If this is so, however, then his strategy for realizing the *concept* orientation requires that he take seriously that *term* orientation which predominated in early phases of terminology work; phases dominated by the concern with (printed) dictionaries, a concern which (if we understand Cimino's views correctly) the *concept* orientation was designed to do away with.

Concepts understood as sets of synonyms, presumably, ought to be seen as standing in different kinds of meaning-relations: *is narrower in meaning*

than, is wider in meaning than, and so forth. Cimino, however, follows the usage now common in much work on biomedical terminologies in speaking of concepts as being linked together also by *ontological* relations, such as *caused by, site of, or treated with* (Cimino, 1998). As I am sure he would be the first to accept, sets of synonymous terms do not stand to each other in causal, locational, or therapeutic relations. In fact, by allowing the latter it seems that Cimino is embracing elements of an *ontological* view of concepts according to which concepts would be abstractions from entities in reality.

I.2.3. *The Ontological View and the Realist Orientation*

On the ontological view, concepts are seen as abstractions of kinds or properties in the real world. This view has advantages over the linguistic and psychological views of concepts when it comes to understanding many of the ways the terms in medical terminologies are, in fact, used by clinicians in making diagnoses. Clinicians refer to *objects*, such as blood clots and kidneys; *properties* which these objects have; and the *kinds* which they instantiate. Cimino, himself, tends toward the ontological view occasionally as, for example, when he refers to the concept *diabetes mellitus* becoming ‘associated with a diabetic patient’ (p. 399). Presumably, this association does not come about because the physician has the patient on his left, and the concept on his right, and decides that the two are fitted together to stand in some unspecified association relation. Rather, there is something *about* the patient, something in reality, which the clinician apprehends and which makes it true that this concept can be applied to this case. Fatefully, however, like other proponents of the concept orientation, Cimino does not address the ontological question of *what it is on the side of the patient* which would warrant the assertion that an association of the given sort obtains. In other words, he does not address the issue of what it is in the world to which concepts such as *diabetes, type II diabetes, or endothelial dysfunction* would correspond.

The ontological view provides us with a means to understand how the corresponding terms can be associated directly with corresponding entities in the biomedical domain. It thereby opens up the question as to the purpose of fabricating concepts to stand in as proxies for those entities. Why should terms in terminologies refer *indirectly* to the world, when doctors and biologists are able to talk about the world *directly*? Of course, the original motivation for fabricating the conceptual realm on the part of

those such as Wüster was the belief that it was impossible to refer to the world directly. But this belief was based on a philosophical presupposition (still accepted today by an influential constituency among philosophers) to the effect that we have direct cognitive access only to our thoughts, not to entities in external reality. By contrast, scientists have never stopped referring to entities in the world directly and, on this basis, have succeeded in constructing theories with remarkable explanatory and predictive power which have undergirded remarkable technological and therapeutic advances. This is one major motivation for our promotion of the *realist orientation*, which we advance as a substitute for the concept orientation, not only because it eliminates the unclarities associated with the latter, but also because of its greater affinity with the methods of empirical science.

On the realist orientation, when scientists make successful claims about the types of entities that exist in reality, they are referring to objectively existing entities which realist philosophers call *universals* or *natural kinds*. A universal can be multiply instantiated by, and is known through, the particular objects, processes, and so forth, which instantiate it. For example, the universal *heart* is instantiated by your heart and by the heart of every other vertebrate. Universals reflect the similarities at different levels of generality between the different entities in the reality which surround us; every heart is characterized by certain qualities exemplified by the universal *heart*, every heartbeat is characterized by certain qualities exemplified by the universal *heartbeat*, and so on.

There is another motivation which we take as supporting a realist orientation. The concept orientation assumes that every term used in a terminology corresponds to some concept in reality and such correspondence is guaranteed; it applies as much to concepts such as *unicorn* or *pneumonia in Russian fiction* as to concepts such as *heartbeat* or *glucose*. However, many terms in medical terminologies are not associated with any universal. There are no universals corresponding, for example, to terms from ICD-9-CM such as:

probable suicide
possible tubo-ovarian abscess
gallbladder calculus without mention of cholecystitis
atypical squamous cells of uncertain significance, probably benign.

Such terms do not represent entities in reality as they exist independently of our testing, measuring, and inquiring activities. Rather, as Bodenreider, *et al.* (2004) point out, they have the status of disguised

sentences representing our ways of gaining knowledge of such entities. This distinction, invisible on the concept orientation, is brought into the light by realism. And it is a distinction which will become increasingly important as automatic systems are called upon to process data in the clinical domain.

It is the existence of universals which allows us to describe multiple particulars using one and the same general term and, thus, makes science possible. Science is concerned precisely with what is *general* in reality; it is interested, not in this or that macrophage, but in *macrophages in general*. It is the existence of such universals which makes diagnosis and treatment possible, by enabling uniform diagnostic and treatment methods (and associated clinical guidelines) to be applied to pluralities of patients encountered in different times and places. In what follows, we will show the advantages that a realist orientation has over the concept orientation in the creation and maintenance of terminologies as well as in other areas of knowledge representation.

I.3. *Concepts are Insufficient for All Areas of Knowledge Representation*

I.3.1. *Some Arguments for the Concept Orientation and Realist Responses*

One argument in favor of conceptualism in knowledge representation is what we can call the *argument from intellectual modesty*, which asserts that it is not up to terminology developers to ascertain the truth of whatever theories the terminology is intended to mirror. This is the job of domain experts. Since domain experts themselves often disagree, a terminology should represent no claims as to what the world is like; instead, it should reflect a conglomeration formed out of the concepts used by different experts.

In fact, however, scientists in medical fields (and other fields) accept a large and increasing body of consensus truths about the entities in these domains. Admittedly, many of these truths are of a trivial sort (that mammals have hearts, that organisms are made of cells), but it is precisely such truths which form the core of science-based ontologies. When there are conflicts between one theory or research community and another, these tend to be highly localized, pertaining to specific mechanisms; for example, of drug action or disease development. Furthermore, such areas of research can serve as loci of conflicting beliefs only because the researchers involved share a huge body of common presuppositions.

We can think of no scenario under which it would make sense to postulate special entities called concepts as the entities to which terms subject to scientific dispute would refer. Since for any such term, either the dispute is resolved in its favor, and then it is the corresponding entity in reality that has served as its referent all along; or it is established that the term in question does not designate anything at all, and the term will then, in the course of time, be dropped from the terminology altogether. The problem that arises from the fact that we do not know, at a given stage of scientific inquiry, whether or not a given term has a referent in reality, cannot be solved by providing such terms with guaranteed referents called concepts.

Sometimes the argument from intellectual modesty takes an extreme form, as in the case of those who consider reality itself to be somehow unknowable (as in, ‘we can only ever know our own concepts’). Arguments along these lines, of course, are familiar not only from the Wüsterian tradition, but also from the history of Western philosophy. Stove provides the definitive refutation (Franklin, 2002). Here we need note only that such arguments run counter not just to the successes, but to the very existence, of science and technology as collaborative endeavors.

The second argument in favor of the concept orientation is what we might call the *argument from creativity*. Designer drugs, for example, are conceived, modeled, and described long before they are successfully synthesized, and the plans of pharmaceutical companies may contain putative references to the corresponding chemical universals long before there are instances in reality. But again, such descriptions and plans can be expressed perfectly within terminologies and ontologies conceived as representing only what is real. For descriptions and plans do, after all, exist. On the other hand, it would be an error to include in a scientific ontology of drugs terms referring to pharmaceutical products which do not yet (and may never) exist, solely on the basis of plans and descriptions. Rather, such terms should be included only at the point where the corresponding instances do, indeed, exist in reality.

Third is what we might call the *argument from unicorns*. According to this argument, some of the terms needed in medical terminologies refer to what does not exist. After all, some patients do believe that they have three arms, or that they are being pursued by aliens. But the realist conception is also equipped to handle phenomena such as these. False beliefs and hallucinations are, of course, every bit as real as the patients who experience them. And certainly such beliefs and episodes may involve

concepts (in the proper, psychological sense of this term). But they are not *about* concepts, and they do not have concepts *as their objects*; for their subjects take them to be about entities in external reality instead. Believing in the concept of aliens in pursuit is not nearly as frightening as believing that there are actual aliens. These patients are making an error, whose proper explanation in our patient records does not consist in asserting that the patients in question, in fact, believed in merely the concept of aliens all along. Such an explanation cannot account for the anxious behavior associated with believing in aliens.

Fourth is the *argument from medical history*. The history of medicine is a scientific pursuit; yet it has often used terms such as ‘phlogiston’ which do not refer to universals in reality. But the domain of the history of medicine is precisely constituted of the beliefs, both true and false, of former generations. Thus, it is expected that a term like ‘phlogiston’ should be included in the ontology of this discipline; not, however, as a free-standing term with a concept as its referent. Rather ‘phlogiston’ should occur as a constituent part of terms denoting the corresponding kinds of beliefs (Smith, 2005b).

Fifth is the *argument from syndromes*. The biological and medical domains contain multitudes of entities which do not exist in reality, but which serve nonetheless as convenient abstractions. For example, a syndrome such as congestive heart failure is an abstraction used for the convenience of physicians for the purpose of collecting under one umbrella term certain disparate and unrelated diseases which have common manifestations or symptoms. Such abstractions are, it is held, mere concepts. From a realist perspective, however, syndromes, pathways, genetic networks, and similar phenomena are fully real, though their reality is that of defined (fiat) classes, rather than of universals. That is, they are real in the sense that they belong to real classes which have been defined by human beings for the very purpose of talking about things which we do not yet fully understand. We may say something similar about the many human-dependent expressions like ‘obesity’, ‘hypertension’, or ‘abnormal curvature of spine’. These terms, too, refer to entities in reality, namely to defined classes which rest on what may be changing fiat thresholds established by consensus among physicians.

Sixth is the *argument from error*. Logical conflicts can arise when falsehoods are entered into a clinical record and interpreted as being about real entities. Rector, *et al.* take this to imply that the use of a meta-language should be made compulsory for all statements in the electronic health

record (EHR). The terms in terminologies devised to link up with such EHRs would refer, not to diseases themselves, but rather merely to the concepts of diseases on the part of clinicians. Thus what is recorded should not be seen as pertaining to real entities at all, but rather to what are called findings (Rector, 1991). Instead of recording both p and $not\ p$, the record would contain entries like: *McX observed p while O'W observed $not\ p$* . Since these entries are about observations, logical contradictions are avoided.

We do not, of course, dispute the fact that clinicians have a perfectly legitimate need to record findings such as an absent finger or an absent nipple. What is disputed, however, is Rector's inference from the fact that there might be falsehoods among the totality of assertions about a given clinical case (or scientific domain), to the conclusion that clinicians (or scientists) should cease to make assertions about the world and, rather, confine themselves to assertions about beliefs.

This proposal contributes to a blurring of the distinction between entities in reality and associated findings. Information about beliefs is fundamentally different in nature from information about objects. Failing to make this explicit allows terminologies to include findings-related expressions in the same category as expressions which designate entities in reality as, for example, in the following assertions from SNOMED CT: 'Genus *Mycoplasma* (organism) *is_a* Prokaryote-cell wall absent (organism) *is_a* bacteria (organism)' and 'Human leukocyte antigen (HLA) antigen absent (substance) *is_a* Human leukocyte antigen (HLA) antigen (substance)'. This running together of two fundamentally different types of assertions introduces obstacles to the working of automatic reasoning systems that employ them as basis.

Of course, we do not deny that clinicians face the need to record, not only the entities on the side of the patient, but also their own beliefs and observations about these entities. Indeed, Rector's argument for the move to conceiving the record as being a record of facts about beliefs rather than of facts about the world is importantly buttressed by appeal to legal considerations which require that the EHR provide an audit trail relating, precisely, to beliefs and actions on the side of medical practitioners. The EHR must serve forensic purposes. From the realist point of view, however, these forensic purposes can be served equally well by a record of facts about the world, as long as we ensure that (a) such facts include facts about beliefs and actions of practitioners (conceived as full-fledged denizens of reality), and (b) the record also preserves data about who

recorded those facts, at what time they were recorded, and so forth, as according to the strategy we outlined in Ceusters and Smith (2006).

On behalf the realist orientation, it can be argued further that even the move to assertions about beliefs would not, in fact, *solve* the mentioned problems of error, logical contradiction, and legal liability. For the very same problems of inadequacy can arise, not only when human beings are describing fractures, pulse rates, coughing, or swellings, but also when they are describing what clinicians have heard, seen, thought, and done. In this respect, these two sets of descriptions are in the same boat, as each is a case of humans describing something. Hence, both are subject to error, fraud, and disagreement in interpretation. The alternative to the Rector approach, we believe, is to provide facilities with the ability to quarantine erroneous entries – and to resolve the concomitant logical conflicts – as they are identified; for example, by appealing to the resources provided by formal theories of belief revision as outlined in Gärdenfors (2003).

The seventh, and final, argument for the concept orientation as a basis for biomedical terminology development is the *argument from borderline cases*. There is often, it is said, no clear border between those general terms which designate universals in reality and those which merely designate classes defined by human beings to serve some purpose. Certainly there are clear cases on either side; for example, ‘electron’ or ‘cell’, on the one hand, and ‘fall on stairs or ladders in water transport NOS, occupant of small unpowered boat injured’ (Read Codes), on the other. But there are also borderline cases such as ‘alcoholic non-smoker with diabetes’, or ‘age dependent yeast cell size increase’, which might seem to call into question the very basis of the distinction.

We will respond, first, with the general point that arguments from the existence of borderline cases usually have very little force. Borderline cases do not undermine the distinction between the entities on either side. The grey area of twilight does not prevent us from distinguishing day from night. Likewise, we can distinguish the bald from the hairy even though we do not know exactly how many hairs one must lose to traverse the border. As to the specific problem of how to deal with borderline expressions of the sorts mentioned – expressions which seem to lie midway between designating universals and designating mere arbitrary classes – this is, in our view, a problem for empirical science, not for terminology. That is, we believe that the normal processes of scientific advance will bring it about that such borderline terms will undergo a filtering process. This process is based on whether they are needed for purposes of fruitful classifications

(for example, for the expression of scientific laws), or for purposes of arbitrary classification (for example, when describing eligible populations for trials).

One generation of scientists may take a given term to refer to a universal, whereas the next generation may discover a reason to believe that the term does not designate anything at all (for example, ‘caloric’), or recognize that it, in fact, refers ambiguously to several universals which must be carefully distinguished (‘hepatitis’). Thus, representational artifacts such as information systems and textbooks, which form an integral part of the practice of science, must be continually updated in light of such advances. But again, we can think of no circumstance in which updating of the sort in question would signify that caloric *is* a concept, or that some expression, at one or other stage, was being used by scientists with the intention of referring to concepts rather than to entities in reality.

I.3.2. *Concepts are Ethereal*

The problematic features of common uses of the term ‘concept’ are not peculiar to the world of biomedical terminology; indeed, they arise generally in the knowledge-representation literature on semantic networks (for example, see *Sowa*, 1992) and conceptual models (Smith, 2006). Here again, concepts (variously called ‘classes’, ‘entity types’, ‘object types’, though information scientists will disagree as to whether the same thing is being expressed by all of these terms) are called upon to perform, at least, two conflicting roles. On the one hand, inside the computer they are delegated to represent concrete entities and the classes of such entities that exist in reality outside of the computer. For example, some abstract proxy – some ghostly *diabetes counterpart* – is required for this purpose, it is held, because *one cannot get diabetes itself inside the computer*. And the computer could reason about diabetes only by creating such a proxy (so the programmer supposes). On the other hand, concepts are delegated to playing the role of representing, in the computer, the knowledge in the minds of human experts. This knowledge is, then, itself characteristically (and again erroneously, as Putnam (1975) argues) assumed to be identifiable with the meanings of the terms such experts use and, in this way, the painful polysemy of ‘concept’ is inherited by the word ‘knowledge’ and its cognates.

Because concepts are pressed into service to perform these various roles, they acquire certain ethereal qualities. Concepts, then, are triply ethereal,

existing in a different sort of denatured guise in the machine, the human mind, and among the meanings stored in language. Their ethereal nature implies that concepts are not the sort of thing that can be examined or inspected. We know what it means to raise and answer questions about, say, a case of diabetes, or about the disease *diabetes* itself. We can turn towards both of these things by directing our attentions to corresponding entities in the world; we can make *what it is on the side of the patient* the target of our mental acts (that to which these acts are directed). We can concern ourselves with traits of the disease or properties of the patient, and we can weigh the separate views advanced by different observers in light of the degree to which they do justice to these traits. But it seems that we can do none of these things in relation to entities in the realm of concepts. The pertinent literature in philosophy and psychology (Margolis, 1999) suggests that concepts are most properly understood, not as *targets* of our cognitive acts, but rather as their *contents*, as that which determines what the target should be and how, in a given act, it should be represented. If this is so, then our puzzlement in the face of questions as to the nature of concepts is understandable. The concept orientation rests precisely on the tacit assumption that concepts would serve as targets – indeed, as the primary targets of concern in work on terminologies – when, in fact, they serve as *contents*.

I.3.3. *The Realm of Concepts Does Not Exist*

A further illustration of the problems associated with the concept orientation is provided by Campbell (1998), in which Keith Campbell, Diane Oliver, Kent Spackman, and Edward Shortliffe – four distinguished figures in contemporary medical informatics – discuss the relevance of the Unified Medical Language System (UMLS, see National Library of Medicine) to current terminology work.

The UMLS Metathesaurus is a well known resource which gathers terms from different source terminologies into a single compendium, with the goal of creating what it calls unified meaning across terminologies. By this its authors mean, roughly, that it creates a framework of common meanings which can be used to provide access to the plurality of meanings carried by terms in the Metathesaurus which derive from a plurality of source terminologies and, consequently, are associated with a plurality of definitions. The purpose is to ensure that everybody who encounters a medical term in a document can use the UMLS to find out the term's

possible meanings. Here, ‘unifying’ is understood as bringing under one framework.

The problem is that the Metathesaurus attempts to do this by creating unified meanings even for those terms which, as they occur in the respective separate source terminologies, clearly have different extensions in the actual world. For example, it assigns the same concept unique identifier (CUI) to both ‘aspirin’ and ‘Aspergum’. In other words, it treats these two terms as if they would refer to (or express) one and the same concept.

Campbell’s (1998) thesis is that this is allowed because there is a Possible World (the authors cite in this connection the work of Leibniz) in which ‘aspirin’ and ‘Aspergum’ do in fact refer to one and the same thing (p. 426). That is, the authors seem to be pointing out that there are situations in which aspirin and Aspergum can be ingested interchangeably. Of course, as the authors admit:

Many clinicians would not regard different formulations of aspirin... as interchangeable concepts in the prescriptions they write. Although aspirin may be an abstract concept, Ecotrin and Aspergum have specific formulations (extensions) in our corporeal world, and use of those particular formulations is subject to different indications, mechanisms of therapy, and risks to the patient. Clearly then, in at least a pharmacy order-entry system, any extensional relationship that was used to determine allowable substitution of pharmacologic formulations would need to have different relationships (representing a different Possible World), than the one currently embodied within the UMLS. However, for a system primarily concerned with the active ingredients of a drug, such as an allergy or drug interaction application, the Possible World embodied in the UMLS may be optimal. (Campbell, 1998, p. 429)

At this point, two questions arise. First, in what sense does the UMLS actually *unify* the meanings of terms? If it only unifies them for certain specific purposes – say, for example, the purposes of those concerned only with a drug’s active ingredients – then it seems to be *restricting* terms’ meanings, rather than unifying them.

Second, in what sense *is* the world, thus defined, *possible*, given that it would have to be governed by laws of nature different from those in operation here on earth? The answer is that it is possible, at best, as an *artifact*, something artificial, inhabiting the same high-plasticity conceptual realm that is postulated by Wüster and his colleagues, a realm in which aspirin may be an abstract concept. In Campbell (1998), the UMLS is itself correspondingly referred to as an artificial world, as contrasted with our

corporeal world of flesh and blood entities. And the job of this artificial world is asserted to be that of providing ‘a link between the realm in which we live and the symbolic world in which computer programs operate’ (p. 426).

Three worlds have hereby been distinguished:

- (1) the possible (‘artificial’) world which is the UMLS,
- (2) the ‘symbolic world’ in which computer programs operate,
- (3) the ‘corporeal world’ in which we live.

How can world (1) link worlds (2) and (3) together? The answer, surely, must involve some appeal to the *extensions* of the concepts in the UMLS. Extensions are understood as collections of the individual objects (actual patients, actual pains in actual heads, actual pieces of Aspergum chewed) in the corporeal world. The authors themselves suggest a reading along these lines when they point out (p. 424), in regard to the terms existing in the UMLS source terminologies, that:

[o]n the one hand there are the physical objects to which [an expression like ‘aspirin’] refers (the expression’s *extensional* component) and on the other there are the characteristic features of the physical object used to identify it (the expression’s *intensional* component).

When it comes to the UMLS itself, however, they abandon this traditional philosophical view in favor of a view according to which (if we have understood their formulations correctly) the extensions of the concepts in the UMLS would be *sets of concepts drawn from source terminologies*:

the developers [of the UMLS] collected the language that others had codified into terminological systems, provided a framework where the intension (connotation) of terms of those systems could be preserved, and unified those systems [into one *unified* system] by providing a representation of extensional meaning by collecting abstract concepts into sets that can be interpreted to represent their extension (p. 425).

They then assert that:

[t]hese extensional sets are codified by the *Concept Unique Identifier* (CUI) in the UMLS. We argue that the ‘meaning’ of this identifier is only understandable extensionally, by examining the characteristics shared by all abstract concepts linked by a CUI (p. 426).

By interpreting ‘extension’ in Wüsterian fashion (which means conceiving extensions in abstraction from the corresponding instances in reality), our authors deny the possibility that the UMLS provides the desired link between the symbolic dimension of computer programs and the domain of real-world entities.

In hindsight, we can see that, with their talk of the UMLS as building a bridge between computers and corporeal reality, Campbell, Oliver, Spackman, and Shortliffe have projected onto the UMLS a goal more ambitious than that which it was really intended to serve. Its actual goal was that of finding unified meaning across terminologies. This weaker goal has proved unrealizable, for the same reason that the concept orientation in general is unrealizable (though there may be some practical value in the imperfect realization even of the weaker goal of unified meaning; for example, in expanding the number of synonyms that can be used to find a target term in a specific terminology). We are still free, however, to readdress the more ambitious goal of building a bridge between computers and corporeal reality, a goal which, with the ineluctable expansion in the use of computers in clinical care (and especially in evidenced-based medicine), becomes ever more urgent.

Part II: *Bridging Computers and the World*

We have claimed that the concept orientation places severe limitations on terminologies to fulfill their potential to support computer applications, a task for which we have claimed that the realist orientation is better suited. In what follows, the reasons for this should become clear.

II.1. *How Terms are Introduced into the Language of Medicine*

Consider what happens when a new disorder first begins to make itself manifest. Slowly, through the official and unofficial cooperation of physicians, patients, public health authorities, and other involved parties, a view becomes established that a certain family of cases, manifesting a newly apparent constellation of symptoms, represents instances of a hitherto unrecognized kind. This kind is a part of reality and, as we have seen above, it corresponds to what realist philosophers call a universal.

The problem is that, in many cases, it is difficult to grasp what universal given particulars are instances of. When a disease universal first begins to make itself manifest in a family of clinical cases, it will be barely

understood. Something similar applies when a new kind of virus or gene, or a new kind of biochemical reaction in the cell, is first detected.

In such cases, a new term is needed to refer to the newly apparent kind. Eventually those involved come to an agreement to use, from here on, (1) *this term* for (2) *these instances* of (3) *this kind*. The concept orientation, however, postulates (4) a new *concept*, together with (5) a *definition*.

II.2. Definitions

On the original ISO-Wüster paradigm, a concept is given what Wüster calls an *intensional definition*, which is an attempt to describe a type of object by referring to characteristic features that its instances have in common. This account works well enough in the relatively straightforward area of woodworking equipment, where Wüster came up with his ideas on concepts and definitions. It works well, too, in a domain like chemistry, for many molecular structures can indeed be precisely and unproblematically defined in terms of exactly repeatable patterns.

However, it confronts two problems in the domain of medicine. One problem occurs in cases where a new universal has only begun to make itself manifest, such as SARS, and it is not yet certain how it is instantiated. Another is that, even if a universal is fairly well understood, we may encounter many instances of it which do not have very many characteristics in common. For example, consider a particular butterfly which might be known to several people, but only at distinct phases of its development. A similar problem is faced when drawing together knowledge concerning successive phases in the development of what is not yet recognized as one single disease.

While in regard to an individual case, users of the term may know precisely what they are referring to (they can point to it in the lab or clinic); nevertheless, it may be difficult to convey this information to others. In such cases, the user has a clear understanding of what the term *designates in reality*, but only at the level of *instances* and not yet at the level of *universals*. As in the case of SARS, or Legionnaires Disease, a term may be introduced as a provisional aid to communication even though the phenomenon has not yet been identified or clearly understood on the level of universals and, on the concept orientation, this means that a new *concept* is thereby introduced in tandem with this term.

There are three strategies which terminologies often employ with respect to providing definitions for new, or problematic, concepts. One is to leave

them undefined, as in the terminology found in SNOMED CT (Bodenreider, 2004). This strategy is itself problematic, for the fewer defined terms a terminology contains, of course, the less value it provides to its users.

The second strategy is for terminologies to fabricate definitions effectively by permuting the constituent words of the term in question. This occurs, for example, in the National Cancer Institute Thesaurus's definition of '*cancer death rates*' as '*mortality due to cancer*'. This practice does not define a term; rather, it merely offers a rewritten version of the term itself. This is akin to defining 'SARS' as meaning *severe acute respiratory syndrome*. This is unhelpful, because not every case of severe acute respiratory syndrome is in fact a case of SARS. The latter covers only those cases of the *severe acute respiratory syndrome first identified in Guangdong, China in February 2002 and caused by instances of a certain particular coronavirus whose genome was first sequenced in Canada in April 2003*. (www.cdc.gov/mmwr/preview/mmwrhtml/mm5217a5.htm)

On the realist orientation, it is recognized that, when more is learned about the new kind that has been discovered, the meaning of the term used to designate that kind will change accordingly. The realist's goal is for a definition to track the development of our scientific knowledge about the world and, ultimately, to capture reality as it is in itself. In our present case, this means capturing that which all instances of a given disease share in common. A real definition provides necessary and sufficient conditions under which it is appropriate to use the term in question as, for example, in this definition taken from the Foundational Model of Anatomy Ontology:

x is a cell =def. *x* is an anatomical structure which has as its boundary the external surface of a maximally connected plasma membrane.

Such a definition describes the real-world conditions under which it is appropriate to use the corresponding term. For many medical terms, only some small number of necessary conditions has been identified thus far. In such cases, it is the job of the definition to describe a partial and still amendable view of what a term actually refers to according to current usage, to be amended as knowledge about it increases.

II.3. *Putting Realism to Work*

Realism sees each terminology as a work in progress reflecting the secure, yet fallible beliefs held at the pertinent stage in the development of

biomedicine about how particular entities in reality are to be classified as instances of universals. Ideally, the result of these works in progress is the increase in the total sum of true beliefs about universals as well as about particulars so that, for example, in biomedicine there is a broad accumulation of knowledge. It is this ideal which the Open Biomedical Ontologies (OBO) Foundry is currently attempting to realize in practice.

Mixed in with such knowledge, however, there will be a small and ever-changing admixture of false beliefs and confusions at every stage. Here, the part of this admixture which concerns us takes the form of terms in a terminology that are believed to refer to some corresponding universal, but which actually do not do so. This can be either because there is no universal at all which can serve as referent of the term in question, or because the term refers ambiguously to what is, in fact, a plurality of universals. With this in mind, we have developed realist counterparts of the three central Cimino desiderata:

Each preferred term in a terminology must correspond to at least one universal (non-vagueness).

Each term must correspond to no more than one universal (non-ambiguity).

Each universal must itself correspond to no more than one term (non-redundancy).

These desiderata are not realizable by any terminological adjustments that are motivated merely by considerations of meaning and language. Rather, they need to be accepted as long-term goals, to which terminologists will come ever closer but never completely realize. In moving towards their realization, terminologists must always follow on the coat-tails of those engaged in empirical research in the attempt to expand our body of knowledge of biomedical universals and their instantiations.

II.3.1. *Knowledge of Universals vs. Knowledge of Instances*

The realist proposal, here, amounts to turning the concept approach onto its head. Whereas the concept approach starts from the top down, letting our thoughts frame our beliefs about reality, the realist approach starts from the bottom up, with the goal of allowing reality itself to form our beliefs about its denizens in a direct way.

Whereas the concept approach admits of only one type of knowledge

(knowledge, precisely, of concepts), the realist approach allows us to distinguish two types: knowledge of universals and knowledge of instances. Knowledge of universals is the sort of *general* knowledge that is recorded, for example, in the textbooks of biomedical science; it is knowledge about the types of entities (such as tuberculosis) that there are in the world. Knowledge of instances is the *particular* knowledge of specific, concrete things (such as this or that particular case of tuberculosis).

We have already seen that it is general knowledge that terminologies are intended to capture, if they are to achieve their practical effect. The domain covered by each terminology comprehends a wide variety of different kinds or categories of universals. In the realm of disorders, these include symptoms, pathological, and non-pathological anatomical structures, acts of human beings (for example anesthetizings, observings), biological processes (disease pathways, processes of development and growth), and more.

In contrast to what is the case in many areas of science, in the domain of clinical medicine knowledge of instances of such universals is of considerable value as well. It is such knowledge that is recorded in clinical records; for example, of patient visits, of emergency call centers, of laboratory results, and so forth. This sort of knowledge is also recorded in automated EHR systems, whose goal is to facilitate clinical data entry in such a way as to enable it to be both used by a human being, and interpreted by a computer application.

The knowledge represented in EHRs is intimately related to the knowledge represented in terminologies. It is through increased discoveries about the sorts of particulars described by EHRs that we gain knowledge about the universals catalogued in clinical terminologies. Obtaining knowledge of a universal, in turn, puts us in a position to recognize particular instances when we come across them. In fact, both kinds of knowledge are indispensable, not only to clinical diagnosis, but to all forms of scientific research.

The better our systems are for keeping track of particulars in the clinical domain, the more efficiently our knowledge of the universals in this domain will be able to advance. However, current EHR regimes embody certain impediments to this advance which, we believe, can be overcome with the help of the realist approach.

II.2.2. *Realism and EHR Systems*

Most existing EHR systems allow direct reference only to two sorts of particulars in reality, namely, (i) *human beings* (patients, care-providers, family members), via proper names or via alphanumeric patient IDs, and (ii) *times* at which actions are performed or observations are made (Ceusters, 2005).

This impoverished repertoire of types of direct reference to particulars means that no adequate means is available to keep track of instantiations of other types of universals (for example, a specific wound, or fracture, or tumor) over an extended period of time. When interpreting health record data, it is correspondingly difficult to distinguish clearly between multiple examples of the same particular, such as *this* tumor, and multiple particulars of the same general kind, such as any tumor existing in patient Brown (Ceusters, 2006).

When a clinician needs to record information about some particular within different contexts – for example, as it exists at different points in time – he must create an entirely new record for each such reference. This is done via some combination of general terms (or associated codes) with designators for particular patients and times; for example, in expressions like, *the fever of patient #1001 observed by physician #4001 at time #9001*. Unfortunately, such composites, even where they are formulated by the same physician using the same general terms deriving from the same coding system, constitute barriers to reasoning about the corresponding entities in software systems, above all because it cannot be unproblematically inferred when such an expression refers to the same entity as does some other, similarly constituted expression. (Imagine a regime for reasoning about human beings as they change and develop over time in which people could be referred to only by means of expressions like, ‘*patient in third bed from left*’, or ‘*person discharged after appendectomy*’, or ‘*relative of probable smoker*’) These sorts of limitations to the knowledge-gathering potential of current EHRs place obstacles in the way of our drawing inferences – for example, for scientific research or public health purposes – from our knowledge of different instances of the same clinical universal in different patients (Ceusters, 2006).

Hence, a way to make the corresponding instances directly visible to reasoning systems is needed (which means visible without need for prior processing). We need to create a regime in which every real-world entity that becomes relevant to the treatment of a patient is explicitly recorded in

the course of data entry. The first step is to expand the repertoire of universals recognized by EHR systems in such a way as to include, in addition to *patient* and *time*, a wide array of other diagnostically salient categories such as *disorder*, *symptom*, *pharmaceutical substance*, *event* (for example an accident in which the patient was involved), *image*, *observation*, *drug interaction*, and so on. When this is done, each entity that is relevant to the diagnostic process in a given case should be assigned an explicit alphanumeric ID – what we have elsewhere called an instance unique identifier (IUI) (Ceusters, 2006) – that is analogous to a proper name.

This would allow EHR systems to do justice to what it is on the side of the patient, in all its richness and complexity. It would also provide an easy means of doing justice to the different views of one and the same instance of a given disorder that may become incorporated into the record; for example, when physician A writes ‘tumor’ and physician B writes ‘CAA12’. The use of IUIs allows us to map the corresponding particulars in our computer representations with one another in a way which serves to make it clear when different physicians are referring to one and the same particular. Indeed, the cumulative result of such use can be understood as a map of the domain in question, showing the multifarious ways in which the universals in the domain relate to one another. In the next chapter, we will see how such maps can be put to use for the purpose of increasing our scientific knowledge of universals and instances alike.

Conclusion

The original motivation for the concept orientation was that it provides a means of representing information which is immune to the vagaries of thought as expressed in natural language. We have shown that the concept orientation, even when Cimino’s desiderata are realized, is beset with flaws which hamper our ability to use terminologies and electronic health records to their full potential. We have advocated a realist orientation, which enables us to bypass the postulation of a conceptual realm and, instead, to engage in the creation of ever-more detailed maps of that reality which, in science and in clinical care, should always be our primary focus.

Chapter 5: The Benefits of Realism: *A Realist Logic with Applications*

Barry Smith

One major obstacle to realizing the general goal of building a bridge between computers and reality on the side of the patient is the existence of multiple, mutually incompatible – and, often impoverished – logical resources bequeathed to those working to improve Electronic Health Record (EHR) systems. In what follows, we will describe a logical framework that is more suitable for the purposes of the realist orientation and provide some examples of how it can be put to use.

1. *The Background of First-Order Logic (FOL)*

In 1879, Gottlob Frege invented the first logical system with a logically perfected language as well as a system of grammatical transformations of the sentences in that language which facilitate processing of information expressed with the language. This system developed into the standard in contemporary symbolic logic, which is known as *first-order logic* (FOL). Contemporary computer languages, such as the Ontology Web Language (OWL), are fragments of FOL which have certain desired computational properties. The language of FOL consists of *individual terms* (constants and variables), representing things in reality; *predicates*, representing properties and relations; *logical connectives* such as ‘and’ and ‘if...then...’; and *quantifiers* (‘for every’, ‘there is some’). The range of variables is normally specified in advance, for example as *all individuals*, *all persons*, *all numbers*, and so forth. The quantifiers are then interpreted accordingly. In some cases quantification is said to be *universal*, and then the range of variables does not need to be specified – it comprehends, in a sense to be specified below, everything.

As an illustration of the use of these ingredients, consider the assertion ‘All horses’ heads are animal heads. In FOL, this would read:

For every individual x , if $\text{horse_head}(x)$, then there is some individual y , such that $\text{animal}(y)$ and $\text{head_of}(x, y)$

Or, to incorporate more of the standard FOL syntax,

$$\forall x[\text{horse_head}(x) \rightarrow \exists y(\text{animal}(y) \ \& \ \text{head_of}(x, y))]$$

Here, the range of variables is all individuals; ‘horse_head’ and ‘animal’ are predicates applied to single individuals; and ‘head_of’ represents a relation between two individuals. To assert that Secretariat is a horse and has a head, we would write:

$$\text{horse}(\text{Secretariat}) \ \& \ \exists x[\text{head}(x) \ \& \ \text{head_of}(x, \text{Secretariat})]$$

treating ‘Secretariat’ as a constant term. To assert that *some* horse has a head, we would write:

$$\exists y [\text{horse}(y) \ \& \ \exists x(\text{head}(x) \ \& \ \text{head_of}(x, y))]$$

First-order logic gets its name because the sentences in first-order language allow quantification (use of ‘for every’, and ‘there is some’) only in relation to what we can think of as ‘first order entities’, which means: entities in the range of the variables (which together form the *universe of discourse*), and thus not in relation to higher-order entities, such as the properties and relations to which the predicates in the language of FOL (‘horse(α)’, etc.) correspond. On standard readings of FOL, the universe of discourse consists only of particular items such as persons or numbers. On these standard readings, to say that quantification is *universal* is to say that when we say ‘for all x , such and such holds’ then we are making an assertion about all *individual entities* in the universe. To make a general statement about objects of a given sort, this statement must be parsed as a conditional assertion. To express the fact that all dogs are four-legged, one has to write a sentence like, ‘for every individual x , if dog(x) then four-legged(x)’. The reader should notice that, given its conditional form (‘if ... then ...’), using this sentence does not commit one to the existence of dogs or of four-legged beings.

FOL’s use of variables hereby allows one to forget that there are real, fundamental distinctions between the sorts of things that exist in reality. In fact, statements about dogs formulated in FOL can be perfectly well conceived as statements about any object in the universe whatsoever, namely that *if* it is a dog, *then* it is four-legged. Here the object plays no

essential role in the sentence. We do not even know, from the standard first-order sentence, whether or not any dogs exist.

2. *A Realist Understanding of First-Order Logic*

In principle, the variables of FOL can range over entities of any sort. In standard practice, however, they have been largely conceived as ranging over individuals (particulars existing in space and time). In keeping with a broadly nominalist slant of most logically orientated philosophers of the 20th century, the universe, from this standard point of view, is the universe of individual things.

In Smith (2005a), an alternative conception of FOL was advanced, differing from this standard conception only in that it deviates, explicitly, from the standard nominalist reading of the range of variables of the original FOL. The alternative view is in the spirit, rather, of Aristotelian realism and accepts, in addition to individual things, universals (kinds, types) as entities in reality. The range of the variables, then, is conceived of as embracing, not only particulars, but also universals. The result is still FOL, in the sense that a distinction is drawn between *predicate* expressions, on one hand, and variable and constant *terms*, on the other. Quantification is still not allowed in relation to the former, and so the logic is still FOL of the perfectly standard sort. But because universals are included in the range of variables, we can now formulate assertions like, ‘there is some quality which John has in virtue of which he is undergoing a rise in temperature’ in this fashion:

For some x [(quality(x) & inheres_in(x , John) & $\exists y$ (rise-in-temperature(y) & causes(x , y)))]

A realist logic of this sort provides the tools needed to deal, in a rigorous way, with real-world instances, and to relate such instances to universals as well as to the general terms used in terminologies. Similarly, drawing on certain ideas worked out in Davidson (1980), it can relate individual things to the processes (events, occurrents) in which they participate.

We can connect general terms to reality by defining the relationships between terms that refer to universals by way of the relationships between their instances (Smith, *et al.*, 2004). In this way, we can provide a simple rigorous account of the relations captured by ontologies such as the Gene

Ontology. Thus, for two universals A and B we can define ‘ A *part_of* B ’ or ‘ B *has_part* A ’ as, respectively:

Every instance of A is part of some instance of B ,

or

Every instance of B has some instance of A as part,

or in symbols:

$$A \text{ part_of } B = \text{def. } \forall x [\mathbf{inst}(x, A) \rightarrow \exists y (\mathbf{inst}(y, B) \ \& \ x \ \mathbf{part_of} \ y)].$$

In other words, A *part_of* B holds if and only if: For every individual x , if x instantiates A then there is some individual y such that y instantiates B and x is a part of y . Correspondingly,

$$B \text{ has_part } A = \text{def. } \forall y [\mathbf{inst}(y, B) \rightarrow \exists x (\mathbf{inst}(x, A) \ \& \ x \ \mathbf{part_of} \ y)],$$

or in other words, B *has_part* A holds if and only if: for every individual y , if y instantiates B then there is some individual x such that x instantiates A and x is a part of y . Here ‘**inst**’ stands for the relation of instantiation between some individual entity and some universal; for example, between Mary and the universal *human being*.

The parthood relations between universals treated by ontologists, hereby, are connected to the more primitive relation of **part_of** between instances, and this is involved, for example, when we say that ‘this finger is part of this hand’, or ‘that step is part of that walk’. Note that assertions using ‘*part_of*’ and ‘*has_part*’ are logically distinct. We can see this, for example, if we consider that, for $A = \textit{cell nucleus}$ and $B = \textit{cell}$, the first is true, but the second is false.

Along the same lines we can define also the ontologist’s *is_a* (is a subtype of) relation as follows:

$$A \text{ is_a } B = \text{def. } \forall x [\mathbf{inst}(x, A) \rightarrow \mathbf{inst}(x, B)].$$

In other words, every instance of universal A is an instance of universal B (as in: all human beings are mammals). We can quantify, too, over universals, for instance if we assert:

$$\forall x [\text{occurrent}(x) \rightarrow \exists y (y \text{ is_a } \textit{continuant} \ \& \ \exists z \text{ inst}(z, y) \ \& \ (z \text{ participates_in } x))]$$

This asserts that, for every occurrent instance, there is some entity y (a universal) which is a subtype of the universal *continuant*, and which is such that at least one of its instances z is a participant in the occurrent x . This ability to quantify over real-world universals and instances is one feature of realist logic that makes it suitable for use in ontology-based information systems. Its flexibility of quantification enables it to be used to track particular instances in EHRs and to link them to universals and, as we will see in Chapter 10, to build terminologies in such a way that their definitions reflect the knowledge that scientists actually have about a given universal, rather than about some associated concepts in their minds.

3. *Concept Logic*

In the 1930s, the great Austrian terminologist Eugen Wüster laid down the central principles of the standard for terminologies propagated by the International Organization for Standardization (ISO) ever since. Unfortunately, instead of adopting FOL, Wüster opted for an older (and weaker) form of concept logic propagated *inter alia* by Kant, in which real-world objects play no essential role.

First-order logic relates each term to instances in reality, and the logic is applied through the process of quantification, which draws the range of its variables from entities in reality. By contrast, instead of relations between terms and entities in reality, CL deals with relations between concepts, such as the *narrower_than* relation, which holds when one concept (for example, *cervical cancer*) is narrower in meaning than another concept (for example, *cancer*). (Thus, CL deals with general terms in the manner of the dictionary maker.) Now, clearly there are a number of connections between this *narrower_than* relation and the ontologist's standard *is_a*. However, from the perspective of CL, *narrower_than* is a relation between meanings which holds, equally, as a relation between mythical or fictional entities as between the entities in reality with which science deals. And, this is the case, similarly, for the other relations of Wüsterian concept logic. For example, ISO (2005) defines the whole-part relation as follows: this relationship covers situations in which one concept is inherently included in another, regardless of the context, so that the terms can be organized into logical hierarchies, with the whole treated as a broader term (p. 49).

Unfortunately, this fixation with concepts results in a logic that is not capable of capturing the logical distinction between universals and instances so that the *part_of* relation between, say, Toronto and Ontario, is treated as identical to that between *brain* and *central nervous system* (see ISO, ‘Guidelines for the Construction, Format, and Management of Monolingual Controlled Vocabularies’, 2005).

A similar concept logic approach underlies much of the work on so-called semantic networks in the AI field in the 1970s (for an overview, see Sowa, 1992). Semantic networks were viewed, initially, with considerable optimism concerning their potential to support what is still called knowledge representation and reasoning (Brachman, 1979). The dawning awareness that this optimism was misplaced was a causal factor in the initial experiments in the direction of what would later come to be called Description Logics (DLs) (Nardi and, Brachman, 2003). The latter fall squarely within the Fregean tradition – effectively, they are a family of computable fragments of FOL – and thus they, too, have some of the resources needed to deal with reasoning about instances. Unfortunately, however, while instances do indeed play a role in the DL world, the instances at issue in DL are often not of this world; thus they are not instances of the sorts encountered, for example, in clinical practice. Work within the DL community – which is often focused on mathematical proxies for real-world instances which exist inside artificial models created *ad hoc* – has led to significant developments in understanding. However, it has served the logicians’ technical purposes of testing consistency and other properties of their systems, rather than the ontologists’ practical purposes of relating a terminology to instances in reality.

With its distinction between T-Box (for terminological knowledge (knowledge about concepts) and A-Box (containing data pertaining to the individual instances in spatio-temporal reality), certainly DL can support reasoning about both concepts and their instances in reality (Brachman, 1979). But the DL community has its roots in the traditional nominalist understanding of FOL, in which the variables and constant terms range over individual things exclusively. Thus, it has paid scant attention to the treatment of instances in different ontological categories; for example, to the differences between instances of attribute kinds (your temperature, your blood pressure) and instances of event or quality kinds (your breathing, your temperature). Similarly, applications of DL-based formalisms in medical terminologies such as GALEN, SNOMED CT, and the National Cancer Institute Thesaurus, have not exploited its resources

for reasoning about instances; rather, they have used the DL-structure as a tool for error-checking on the terminological level. And this is so, in spite of the fact that one central purpose to which such terminologies could be applied is to support the coding of EHRs which relate, precisely, to instances in reality.

4. 'Terminology' Defined

Terminologies have certain parts and structures in common. Delineating these parts and structures will help us to obtain an explicit understanding of what a terminology is and, hence, of the advantages a terminology can provide if it is constructed along the lines of a realist orientation.

In order to understand its components and structure, we may describe a terminology more technically as a graph-theoretic object (of the sort presented in Figure 1) consisting of nodes joined together by links, the whole indexed by version number. Multi-sorted logic enables us to codify this information into a formal definition of 'terminology' (Smith, *et al.*, 2006).

What are the common components of terminologies? First, are nodes, represented as the tips of branch-like structures. There are three kinds of information which a node may contain, namely, (1) a preferred term p , (2) any synonyms S_p which this term may have, and (3) (ideally) a definition d for that term (and its synonyms).

Figure 1: *Graph-theoretic Representation of the FMA Terminology*

The screenshot shows the Foundational Model Explorer (FME) interface. At the top, there is a logo for FME and the title 'Foundational Model Explorer' with 'Options' and 'Help' buttons. Below the title bar, there is a search bar and a dropdown menu for 'Select navigation tree type:' set to 'part'. The main area is divided into two panes. The left pane shows a tree view of the FMA Terminology, with 'Neuraxis' selected and expanded to show its sub-terms: Autonomic nervous system, Somatic nervous system, Enteric nervous system, Peripheral nervous system, Cardiovascular system, Respiratory system, Alimentary system, Urinary system, Female genital system, Endocrine system, Sense organ system, Deep fascial system, and Stomatognathic system. The right pane shows the details for the selected 'Neuraxis' node, including its preferred name, synonyms, and non-English equivalents.

Foundational Model Explorer Options Help

Search

Select navigation tree type: part

+ Neuraxis
 + Autonomic nervous system
 ▪ Somatic nervous system
 ▪ Enteric nervous system
 + Peripheral nervous system
 + Cardiovascular system
 + Respiratory system
 + Alimentary system
 + Urinary system
 + Female genital system
 + Endocrine system
 ▪ Sense organ system
 + Deep fascial system
 ▪ Stomatognathic system

PREFERRED NAME:

SYNONYM:

Central nervous system
Cerebrospinal axis
CNS

NON-ENGLISH EQUIVALENT:

name	language
Pars centralis	Latin
Systema nervosum centrale	Latin
Système nerveux central	French

There are various different ways in which nodes can relate to one another in such a graph; for example, lower nodes can relate to higher ones in relations such as *part_of*, *is_a*, and so forth (for more on relations see Chapters 10 and 11). These relations among nodes are represented by *links* (L), the second kind of information which terminologies contain. Links may be represented visually as the branches which connect the nodes. Reality contains almost an infinite number of relations in which entities may stand to one another. Ideally, there would be as many kinds of links as there are kinds of relations. Realistically, however, a terminology is limited, and can only contain information about the most salient relations obtaining between the entities represented by terms in its nodes. Links contain two kinds of information, namely, (1) a description of the relation itself (*r*), and (2) a description of the way in which the relation obtains between the terms which the link connects (L_r , which describes $p r q$). Of course, these relations must either be explicitly defined or taken as primitives; in the latter case, they must be explicitly axiomatized so that their meaning is made clear.

The third kind of information contained in terminologies pertains to the particular time (*t*) at which a particular version of a given terminology is in use. On a realist, scientifically oriented and evidence-based conception, our terminologies ought to evolve as our knowledge of the world evolves. It is crucial to keep track of these changes in our knowledge so that we know how terms are used now, and of the ways in which terms were used previously for describing our previous working view of what the world was like. Hence, each version of a terminology must be indexed according to a particular time.

We can use a realist logic to provide a precise definition of a terminology and, thereby, to record information about terminologies themselves. Let n_1, n_2, n_3, \dots name individual nodes in a terminology graph. Let L_1, L_2, L_3, \dots name individual links. Let v_1, v_2, v_3, \dots stand in for particular dates.

A terminology, then, is an ordered triple: $T = \langle N, L^*, v_n \rangle$,

where: N is the set of *nodes* n_1, n_2, n_3, \dots in the terminology, where each n_i is a triple $\langle p, S_p, d \rangle$, with p a *preferred term*, S_p a set of *synonyms*, and d a *definition* (ideally). L^* is the set of L_1, L_2, L_3, \dots where each L_i is a *link* that consists of an ordered pair $\langle r, L_r \rangle$, consisting of a relation designation r (*'is_a'*, *'part_of'*, etc.), together with a set L_r or ordered pairs $\langle p, q \rangle$ of

those preferred terms for which ‘ prq ’ represents a consensus assertion of biomedical science about the corresponding universals at the time when the given terminology is prepared, and v_n is a version number, which encodes this time.

On our realist account, the variables $p, q, d, r, v\dots$ stand simply and unambiguously for syntactic entities, or strings of characters in some regimented language. These syntactic entities include what are called preferred terms, which are the officially recommended representations of given universals in reality. Such preferred terms are recorded in the terminology, along with the various synonyms (the ways of referring to this universal) used by sub-communities of specialists. Such preferred terms may prove to be erroneous; that is, we may *discover* through scientific inquiry that a given term (for example ‘phlogiston’, or ‘aura’) corresponds to no universal and, thus, to no instances in reality.

By contrast, according to the concept orientation the mentioned variables are seen as ranging, not over syntactic strings, but over concepts in people’s minds. From the perspective of the concept orientation, there is a one-to-one correspondence between preferred terms and concepts, and this has the unfortunate result that every preferred term in a terminology is guaranteed a referent. So, for example, on the concept orientation there is no way to express the discovery that the term ‘caloric’ does not, in fact, correspond to anything in reality at all.

Our realist account creates no such problem. Some terms within the range of our variables will not correspond to a universal in reality; like ‘unicorn’, ‘phlogiston’, or ‘caloric’, they will be *empty names*. Other terms represented by these variables will have the opposite problem in that they will correspond to *too much* in reality, that is, they will refer ambiguously to a plurality of universals. When evaluating terminologies, we need to take both of these alternatives into account by considering the entire terminology $T = \langle N, L^*, v \rangle$ in light of its status as a map of an analogous structure of universals on the side of reality.

In the ideal situation, where all of our terms perfectly represent universals in reality, we could indeed associate N in one-to-one fashion with some corresponding set U of the universals designated by its constituent nodes. However, really existing terminologies fall short of this ideal in the three ways identified in what we can think of as realist counterparts of Cimino’s criteria of non-vagueness, non-ambiguity, and non-redundancy (Cimino, 1998). This means (roughly, and for our present

purposes) that, at any given stage, the nodes of any terminology will be divided into three groups N_1 , $N_>$, and $N_<$. In other words,

$$N = N_1 \cup N_> \cup N_<$$

where N_1 consists of those in nodes in N whose preferred terms correspond to exactly one universal, $N_>$ of those nodes in N whose preferred terms correspond to more than one universal (in various combinations), and $N_<$ of those nodes in N whose preferred terms correspond to less than one universal (in the simplest case, to no universal at all).

Our realist account assumes that, with the passage of time, $N_>$ and $N_<$ will become ever smaller, so that N_1 will approximate N ever more closely. However, this assumption must be qualified in reflection of the fact that N is itself changing, as our knowledge of the salient universals in biomedical reality expands through new discoveries.

Our knowledge of the successes medical science has had to date gives us strong reason to believe that N_1 constitutes a large portion of N . N , remember, is a collection of terms already in use, each one of which is intended to represent a biomedical universal. N includes very many presently uncontroversial terms which we are normally inclined to overlook, such as ‘heart’ or ‘tumor’. At the same time, our knowledge of the ways errors continue to be uncovered in specific terminologies gives us reason to believe that we have some way to go before $N_>$ and $N_<$ can be excised completely, if this will ever be possible.

Moreover, we know *a priori* that at no stage (prior to that longed-for end to our labors that seems forever just out of reach) will we know precisely where the boundaries are to be drawn between N_1 , $N_>$, and $N_<$, that is, we will never know precisely which portions of N consist of the low value $N_>$ - and $N_<$ -type terms. The reason for this is clear; if we *did* know where such terms were to be found, then we would already have the resources needed to expand the size of N_1 correspondingly and, hence, to move its boundaries to a different position closer to N .

However, on the realist orientation this unavoidable lack of knowledge of the boundaries of N_1 is not a problem; since it is, after all, N , and not N_1 , which is the focus of the practical labors of ontologists. It is N which represents our (putative) consensus knowledge of the universals in the relevant domain of reality, at any given stage. Thus the *whole* of N , as far as the developers and users of a given terminology are concerned, consists of *names of universals*.

But if we do not know how the terms are presently distributed among the three groups, does this mean that the distinction between N_1 , $N_>$, and $N_<$ is of purely theoretical interest, a matter of abstract philosophical housekeeping that is of no concrete significance for the day-to-day work of terminology development and application? Not at all. Typically, we will have, not just one version of a terminology, but a developing series of terminologies at our disposal. In uncovering errors immanent to a terminology, we thereby uncover terms which must be excluded from future versions because they do not correspond to universals. Given the resources of our realist approach, however, we do not need to wait for the actual discovery of error; for we can carry out experiments with terminologies themselves, which means that we can explore through simulations the consequences of different kinds of mismatch between our terms and reality. For more detail see Ceusters (2006), Ceusters and Smith (2006), and Ceusters, Spackman and Smith (2007).

5. A Formal Framework for Terminology Experimentation

Once again, consider our scenario of the way in which a medical term describing a disease or a disorder is introduced into our language. The *instances* in our initial pool of cases, as well as certain regularities and patterns of irregularities (deviations from the norm) which they exemplify, are well known to the physicians involved. However, the *universal* which they instantiate is unknown. The challenge, in this case, is to *solve* for this unknown in a manner that is similar to the way in which astronomers postulated an unknown heavenly body, later identified as Pluto, in order to explain irregularities in the orbits of Uranus and Neptune. Three different kinds of solution can present themselves, as the cases of disorders in the pool are either (i) instances of exactly one universal, (ii) instances of no universal at all, or (iii) instances of more than one universal.

In what follows, we will present a rigorous framework which is designed to put us in a position where we can extract certain kinds of valuable information from the resources provided by terminologies and EHRs. We believe that, in the long run, this information can enable terminologies and EHRs to play much larger roles in making themselves amenable to quality control, supporting decisions in the process of diagnosis of medical disorders, and facilitating scientific discoveries.

Note that this idea will only be realizable in a future world of sophisticated EHRs in which instances in clinically salient categories are

tracked by means of instance unique identifiers (IUIs) of the sort described in Chapter 4. Each such IUI would be associated with other relevant information about the disorder or disease in question as it is expressed in a particular case. We can think of the result as a vector (an ordered n -tuple) of instance-information, comprehending coordinates for the following kinds of information: (1) the relevant terms in one or more terminologies; (2) cross-references to the IUIs assigned to those other particulars (such as patients) with which the disorder under scrutiny is related; and (3) the measured values of relevant attributes such as temperature and blood pressure, as well as bio-assay data such as gene expression. Each coordinate will then be indexed by time of entry, source, and estimated level of evidence.

We will call the sum of all information that is pertinent to a particular manifestation of a disorder an *instance vector*. A definition of ‘instance vector’ will thus include variables for each of the following components: i an IUI, a preferred term p in a terminology, and the designation of a time at which the particular catalogued by i is asserted to be an instance of the universal (if any) designated by p (for details Ceusters and Smith, 2007). Thus, an instance vector can be expressed as an ordered triple, $\langle i, p, t \rangle$.

Suppose, for example, that i corresponds to patient Brown’s hernia, p to the term ‘hernia’, and t to the time at which his hernia was discovered. Our goal is to see formally how a given terminology at a given time is linked to a given set of IUIs (containing information gathered for example by a single healthcare institution during a given period). In order to achieve this, we need a formal way of representing a terminology as it exists at a given time and as it corresponds with a set of instance vectors. We will call this combination of terminological information with instance information a t -*instantiation*, represented by the variable I_t . Thus, for a given set D of IUIs, we can define a t -instantiation $I_t(T, D)$ of a terminology $T = \langle N, L^*, \nu \rangle$ as: the set of all instance vectors $\langle i, p, t \rangle$ for i in D and p in N . For example, each record containing the IUI corresponding to patient Brown’s hernia (i) at time t , where i is a IUI that is a member of the set D and ‘hernia’ is a term (p) in the terminology N .

Next, we need a way to map the extension of the universal designated by the term p in the particular domain of reality selected for by D at time t , assuming that p does indeed designate a universal (we address this assumption below). In other words, we want to define for each term p the set of all IUIs for which the instance vector is included in the t -instantiation. We will call this the t -*extension* of p .

Our definition of t -extension enables us to examine, for each term p , its t -extensions for different values of D and t . This will enable us, in turn, to determine statistical patterns of different sorts, taking into account also, for each i , the other instance vectors in which i is involved through the relations in which the corresponding instances stand to other instances represented by IUIs in D . Our three alternative scenarios will then, once again, present themselves according to the status of each preferred term p in relation to the world of actual cases (the world which serves as standard for the truth and falsity of our assertions):

1. p is in N_1 (there is a single universal designated by p) and, in this case, the instances in $I_t(T, D)(p)$ have in common a specific invariant pattern (which should be detectable through the application of appropriate statistically based tools);
2. p is in $N_>$ (p comprehends a plurality of universals, for example in a manner analogous to the term ‘*diabetes*’) and, in this case, the instances in $I_t(T, D)(p)$ manifest no common pattern, but they (or the bulk of them) can be partitioned into some small number of subsets in such a way that the instances in each subset do instantiate such a pattern;
3. p is in $N_<$ (p corresponds to no universals) and, in this case, the instances in $I_t(T, D)(p)$ manifest no common pattern, and there is no way of partitioning them (or the bulk of them) into a combination of one or a small number of subsets in such a way that all the instances in each subset instantiate such a pattern.

6. Reasoning with Instance Identifiers: Three Applications

There are at least three applications for a system along the lines described. Such a system could be used, first of all, for purposes of quality-control of terminologies (and thus, for purposes of automatically generating improved versions of terminologies). For a given disorder term p , we gauge whether p is in N_1 , $N_>$, or $N_<$ by applying statistical measures to the similarities between the vectors associated with each of the members of relevant instantiations. For example, two vectors are *similar* if the data they contain are close numerically (say, if two times are close to one another in a sequence), or if two terms represent the same or similar types,

or if they represent the same entity on the instance level (say, a set of IUIs signifies the same disorder in the same patient).

Here is an example of the benefit of applying statistical measures to the similarities between vectors. If the measure of similarity between vectors is both roughly similar for all members of a given instantiation and also roughly constant across time when measures are applied to instances for which we have similar amounts of data of similarly high evidence-value, then this will constitute strong evidence for the thesis that p is in N_1 . If, on the other hand, we find high similarity for some disorder term before a certain time t but much lower degrees of similarity after some later time t^+ , then we can hypothesize that the relevant disorder, itself, has undergone some form of mutation, and we can experiment with adding new terms and then repartitioning the available sets of IUIs in such a way as to reach, once again, those high levels of similarity which are associated with the N_1 case.

In due course, such revision of terminologies will give rise, in the opposite direction, to revisions of the information associated as vectors to each of the relevant IUIs. We might, for example, discover that a given single disorder term has thus far been applied incorrectly to what are in fact instances of a plurality of distinct disorders. Such revision will lead, in turn, to better quality clinical record data, which may give rise to further revisions in our terminologies.

Second, such methods for reasoning with terminology and instance data might be used for supporting decisions in the process of diagnosis. In a world of abundant instance data, one goal of an adequate terminology-based reasoning system would be to allow the clinician to experiment with alternative term-assignments to given collections of instance data in ways which would allow measurements which result in the greater and lesser likelihood of given diagnoses, on the basis of statistical properties of the patterns of association between terms and instances. Thus, we could imagine software which would allow experimentation with alternative IUI and term assignments; for example, when it is unclear whether successive clusters of symptoms in a given patient should be counted as manifestations of single or of multiple disorders. The machinery of instantiations, then, could be used to test out alternative hypotheses regarding how to classify given particulars by offering us the facility to experiment with different scenarios as concerns the division between N_1 , $N_<$, and $N_>$ in relation to given cases.

In the real world, of course, such methods cannot be applied successfully in every case. For example, we may not have all the data needed to convince a computer armed with a stock of universal terms and associated instance data that a given case meets the requirements for any available diagnosis. Such a situation, however, is no different from that which is faced already by the practicing physician, who must decide from case to case how much data to collect (for example, how often to take the temperature of a given patient) in order to achieve a succession of better approximations to what then establishes itself as a good diagnosis. He learns how to do this, first, from medical textbooks and education, then through experience and by following guidelines and protocols.

Finally, the methodology presented here can be used to facilitate scientific discoveries. Suppose, for example, that the length of a patient's nose is correlated with a certain specific disease, but that this fact is unknown to medical science. Why should anyone start to register the patient nose-length in the way that we do now for, say, temperature or blood pressure? The answer is that we do so already. Many hundreds of thousands of patients have undergone plastic surgery for cosmetic nose corrections. In each case, the length of the nose is measured as a matter of course. Many of these patients visited other physicians for totally different problems (before, at the same time, or later). If all the physicians involved had been exploiting the potential of referent tracking as here conceived, then it would not be difficult to correlate these data, using brute-force techniques such as cluster analysis, principal component analysis, or factor analysis, to tease out the correlation in question, in just the way that scientific discoveries are sometimes made on the basis of instance-level data in other domains. (For more details see; Ceusters W, Smith B. 'Referent Tracking and its Applications'.)

7. Conclusion

When we take advantage of realist (instead of conceptual) logic, we can harness the information provided by these maps to accelerate our gains in knowledge about the world by keeping track of the instances which fall within the range of the variables of our logic. In the ideal case, a biomedical terminology would provide, not merely the resources for assigning preferred terms for universals to the corresponding instances in reality, but also a perspicuous map of how these universals themselves are related to each other in reality. As we conceive the EHR systems of the

future, instance data will be, to a large degree, automatically partitioned at the point of data entry in ways reflecting the structure of the world of clinically relevant universals. Currently, this partitioning of instances is masked from view in the clinical record because the instance-level data that exists in separate EHRs is accessible only via the detour of reference to the individual patient. A regime for the management of terminologies and clinical data along the lines described above, however, would allow us to directly map the instances that are salient to medical care in such a way as to mirror how the latter are related together in reality at the level of both instances and universals. In this way, it would make a new level of sophistication in reasoning about *what it is on the side of the patient* possible, which is the primary focus of medical care.

Chapter 6: A Theory of Granular Partitions

Thomas Bittner and Barry Smith

1. *Introduction*

Imagine that you are standing on a bridge above a highway checking off the makes and models of the cars that are passing underneath; or a laboratory technician sorting samples of bacteria into species and subspecies; or you are making a list of the fossils in your museum. In each of these cases, you are employing a certain grid of labeled units, and you are recognizing certain objects as being located in those units. Such a grid of labeled units is an example of what we shall call a *granular partition*. We shall argue that granular partitions are involved in all naming, listing, sorting, counting, cataloguing, and mapping activities, activities that are performed by human beings in their traffic with the world. Partitions are the cognitive devices designed and built by human beings to fulfill these various listing, mapping, and classifying purposes.

2. *Types of Granular Partitions*

Some types of granular partition grids are flat and amount to nothing more than mere lists. Others are hierarchical, consisting of units and subunits, the latter being nested within the former. Some grids are built in order to reflect independently existing divisions on the side of objects in the world (such as the subdivision of hadrons into baryons and mesons). Others – for example, the partitions created by classifying organisms into phyla or kingdoms, or by electoral redistricting – are themselves such as to create the corresponding divisions on the side of their objects, and sometimes they create those very objects themselves. Quite different sorts of partitions – having units of different resolutions and effecting unifying, slicings and reapportionings of different types – can be applied to the same domain of objects simultaneously. Members of the animal kingdom can be divided according to what they eat, where they are indigenous, or even by the number of appendages an individual animal has. Maps, too, can impose subdivisions of different types upon the same domain of spatial reality, and the icons which they employ represent objects in granular fashion (which means that they do not represent the corresponding object parts).

The theory of partitions is highly general, and this generality brings a correspondingly highly general reading of the term *object* with it. Here, we take an object to be any portion of reality like an individual, a part of an individual, a class of individuals (for example, a biological species), a spatial region, a political unit (county, polling district, nation), or even the universe as a whole. An object in the partition-theoretic sense is everything existent that can be recognized by some unit of a partition.

Objects can be either of the bona fide or of the fiat sort (Smith, 2001). Bona fide objects are objects which exist (and are demarcated from their surroundings) independently of human partitioning activity. Fiat objects are objects which exist only because of such partitioning activity. In some cases, partition units recognize fiat objects, such as your right arm or Poland, which exist independently of human cognition but which have boundaries that depend upon our human demarcations. In other cases, fiat objects are created through the very projection of partition units onto a corresponding portion of reality. Partitions which themselves create fiat objects include, for example, the partitions created by our ordinary classification of fruits and vegetables. Once fiat objects have been created in this way, subsequent partitions may simply recognize them (without any object-creating effect), as do those partitions which recognize bona fide objects.

In first approximation, granular partitions can be conceived as the mereological sums of their constituent units. This conception is roughly analogous to David Lewis's (1991) conception of classes as the mereological sums of their constituent singletons (thus the set $\{1, 2, 3\}$, for example, is conceived by Lewis as the mereological sum of the three sets $\{1\}$, $\{2\}$, and $\{3\}$). The units within a granular partition, however, may manifest a range of properties which the singletons of set theory lack. This is so because, where each singleton is defined in the obvious way in terms of its single member, a unit of a granular partition is determined by its *label* as well; and this means independently of any object which might fall within it. The units of a partition are what they are independently of whether there are objects located within them. A map of Middle Earth is different from a map of the Kingdom of Zenda, even though, in both cases, there is nothing on the side of reality upon which these maps would be projected.

Just as when we point our telescope in a certain direction we may fail to find what we are looking for, so too, when we point our partition in a certain direction it may be that there are no objects located in its units. But

this does not mean that the theory of partitions recognizes some counterpart of the set theorist's empty set (an entity that is contained as a subset within every set). For the empty set is empty by necessity; by contrast, a unit in a partition is empty *per accidens* at best.


The theory of Granular Partitions was originally developed in terms of first-order predicate logic, and a corresponding presentation of the theory in formal logical notation can be found in Bittner and Smith (2003). In what follows, however, symbolic notation is avoided as much as possible.

3. *Granular Partitions as Systems of Labeled Units*

3.1. *A Bipartite Theory*

In the present essay, we present the basic formal theory of granular partitions. Our formal theory has two orthogonal and independent parts, namely, (A) a theory of the relations between units, subunits, and the partitions in which they are contained, and (B) a theory of the relations between partitions and objects in reality. In a set-theoretic context, the counterpart of (A) would be the study of the relations among subsets of a single set, while the counterpart of (B), in the same context, would be the study of the relations between sets and their members. Partition theory departs from the extensionalism of set theory (that is, from the assumption that each set is defined exclusively by its members, so that two sets are identical if and only if they have the same members). A unit is defined by its position within a partition and by its relations to other units, and it is this which gives rise to the relations treated of by theory (A). What objects in reality are located in a unit – the matter of theory (B) – is a further question, which is answered in different ways from case to case. Briefly, we can think of units as being projected onto objects in something like the way in which flashlight beams are projected upon the objects which fall within their purview. For partition theory seeks to represent the ways in which cognitive agents categorize reality as it presents itself in all its un-mathematical variety and scruffiness.

Consider the left part of Figure 1. Theory A governs the way we organize units into nesting structures and the way we label units. Theory B governs the way these unit-structures project onto reality, as indicated by the arrows connecting the left and the right parts of the figure.

Figure 2: *Relationships between Units and Objects*

3.2. *The Subunit Relation*

Theory (A) is, effectively, a formation theory for partitions; it studies properties partitions have in virtue of the relations between, and the operations performed upon, the units out of which they are built independently of any linkage to reality beyond. Units in partitions may be nested one inside another in the way in which, for example, the species *crow* is nested inside the species *bird* in standard biological taxonomies. When one unit is nested inside another in this way, we say that the former is a *subunit* of the latter.

We use z, z_1, z_2, \dots as variables ranging over units and A, A_1, A_2, \dots as variables ranging over partitions. We write ' $z_1 \subseteq_A z_2$ ' in order to express the fact that z_1 stands in a subunit relation to z_2 within the partition A . (Where confusion will not result, we will drop the explicit reference to the partition A and write simply ' \subseteq '.) We can state the first of several master conditions on all partitions as follows:

MA1: The subunit relation \subseteq is reflexive, anti-symmetric, and transitive.

This means that within every partition: each unit is a subunit of itself; if two units are subunits of each other, then they are identical; and if unit z_1 is a subunit of z_2 and z_2 a sub-unit of z_3 , then z_1 , in its turn, is a subunit of z_3 .

3.3. *Existence of a Maximal Unit*

A maximal unit is the unit in a partition which encompasses all of the other subunits. (However, as we will see below, the maximal unit is *not* identical to the largest partition.)

DMax: A unit z_1 of partition A is a *maximal unit* if and only if every unit of A is a subunit of z_1 .

We now demand, as a further master condition, that

MA2: Every partition has a unique maximal unit in the sense of DMax.

The motivation for MA2 is very simple, and turns on the fact that a partition with two maximal units would either be in need of completion by some extra unit representing the result of combining these two maximal units together into some larger whole, or it would be two separate partitions, each of which would need to be treated in its own right.

MA2 implies that there are no partitions which are empty *tout court* in that they have no units at all. Although the maximal unit seems to be just as big as the partition itself, it is not identical to it; for the maximal unit is just that, a unit. Further, it is one that comprehends only a very large-grained perspective. By contrast the partition, itself, includes all the other subunits as its units as well.

3.4. *Finite Chain Condition*

The transitivity of the subunit relation generates a nestedness of units inside a partition in the form of chains of units, structured in such a way that maximal unit z_1 contains z_2 , which contains z_3 , all the way down to some z_n . We shall call the units at the lower ends of such chains *minimal units* (also called leaf nodes in the terminology of ontologies) and define:

DMin: z_1 is a *minimal unit* of partition A if and only if every subunit z of partition A which is a subunit of z_1 is also identical to z_1 .

Another important aspect of a partition, then, is:

MA3: Each unit in a partition is connected to the maximal unit via a finite chain.

MA3 does not rule out the possibility that a given unit within a partition might have infinitely many immediate subunits, those which have no other units as intermediaries between themselves and their parent units. Thus MA3 enforces finite chains between the maximal partition and each of its units: every sequence of descending sub-units stops after a finite number of

steps and there is only a finite number of such sequences; but this leaves open the issue of whether partitions themselves are finite.

If, in counting off the animals you saw in the rainforest, your checklist includes one unit labelled *located in trees* and another unit labelled *has fur*, we will rightly feel that there is something amiss with your partition. One problem is that you will almost certainly be guilty of double counting. Another problem is that there is no natural relationship between these two units, which seem rather to belong to distinct partitions. As a step towards rectifying such problems we shall insist that all partitions must satisfy the condition that every pair of distinct units within a partition are either related by the subunit relation or are disjoint. In other words:

MA4: If two units within a partition overlap, then one is a subunit of the other.

From MA3 and MA4 we can prove, by a simple *reductio*, that the chain connecting each unit of a partition to the maximal unit is unique.

3.5. *Partition-theoretic Sum and Product of Cells*

In this chapter, the background to all our remarks is mereology, which is the study of the relationships between wholes and their parts. We take the relation \leq meaning ‘part of’ as primitive, and define the relation of overlap between two entities, simply, as the sharing of some common part. The part-of relation is like the subunit relation in being reflexive, anti-symmetric, and transitive, but the two differ in the fact that the subunit relation is a very special case of the parthood relation.

The subunits of a unit are also parts of the unit just as, for David Lewis (1991), each singleton is a part of all the sets in which it is included. What happens when we take the mereological products and sums of units existing within a partition? In regard to the mereological product of two units, $z_1 * z_2$, matters are rather simple. This product exists only when the units overlap mereologically, that is, only when they have at least one subunit in common. This means that the mereological product or intersection of two units, if it exists, is in every case just the smaller of the two units.

In regard to the mereological *sum* of units $z_1 + z_2$, in contrast, a more difficult situation confronts us. Given any pair of units within a given partition, there is a corresponding mereological sum simply in virtue of the fact that the axioms of mereology allow unrestricted sum-formation.

However, this mereological sum will be a unit within the partition in question only in special cases. This occurs, for example (and simplifying somewhat), when units labeled *male rabbit* and *female rabbit* within a partition have the unit labeled *rabbit* as their sum. By contrast, there is no unit labeled *rabbits and jellyfish* in our standard biological partition of the animal kingdom.

To make sense of these matters, we need to distinguish the mereological sum of two units from what we might call their *partition-theoretic sum*. Their mereological sum is the result of taking the two units together in our thoughts and treating the result as a whole, while their partition-theoretic sum consists of those mereological sums which we can recognize against the background of a given partition. The partition-theoretic sum of units z_1 and z_2 in a partition is the smallest subunit within the partition containing both z_1 and z_2 as subunits; that is, it is the least upper bound of z_1 and z_2 with respect to the subunit relation. (By MA2 and MA4, we know that this is always defined and that it is unique.) In general, this partition-theoretic sum is distinct from the mereological sum of the corresponding units. (The partition-theoretic sum of the units labelled *rabbit* and *lion* is the unit labelled *mammal* in our partition of the animal kingdom.) The best we can say, in general, is that the mereological sum of z_1 and z_2 ($z_1 + z_2$) is at least part of their partition-theoretic sum (or $z_1 \cup z_2$) (Smith, 1991). On the other hand, note that if we analogously define the partition-theoretic product ($z = z_1 \cap z_2$) of two units within a given partition as the largest subunit shared in common by z_1 and z_2 – that is, as their greatest lower bound with respect to the subunit relation – then it turns out that this coincides with the mereological product already defined above. Mereological sum and product apply to both units and objects, while partition-theoretic sum applies to units only. The following symbols are used for the two groups of relations:

Figure 2: *Symbols for Relations of Partition-theory and Mereology*

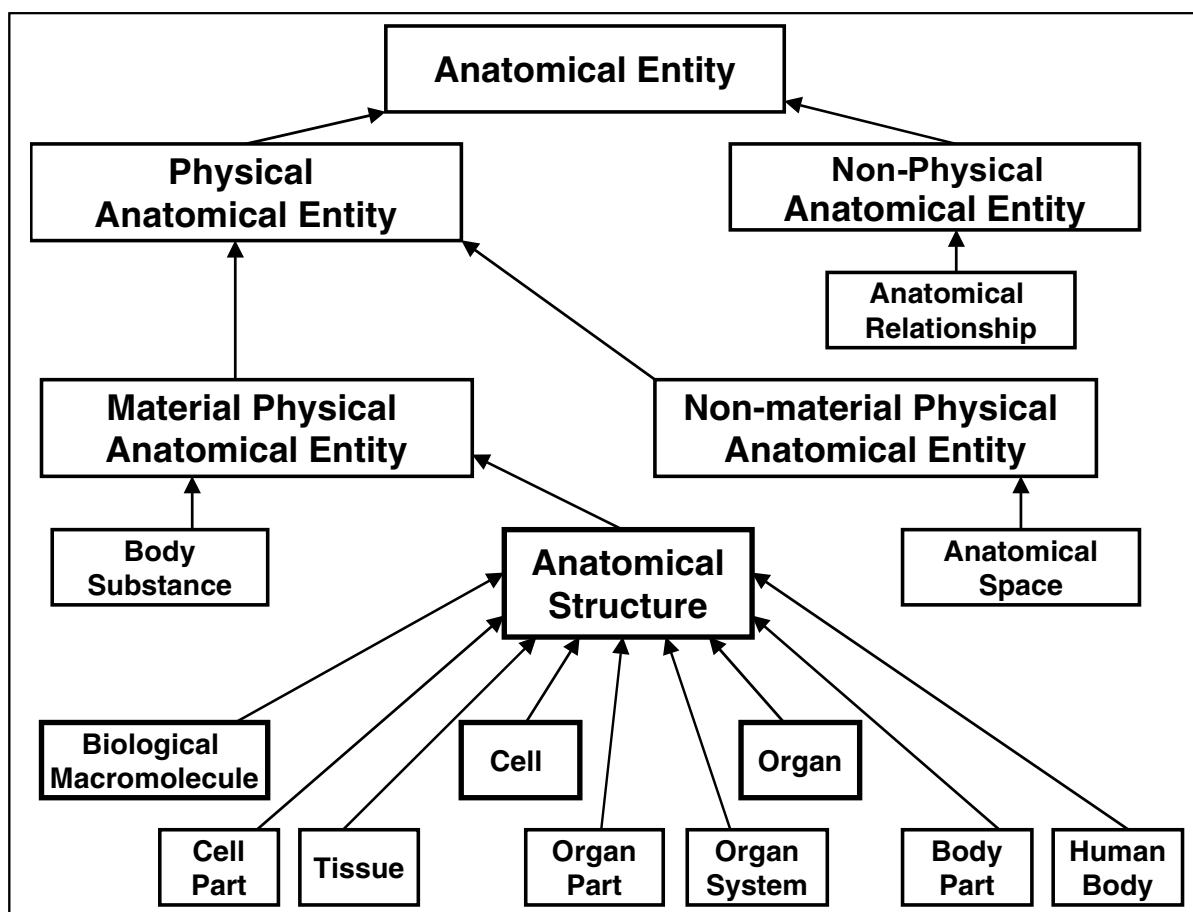
	Partition-theoretic (for units)	Mereological (for units and for objects)
Sum	\cup	+
Product	\cap	*
Inclusion	\subseteq	\leq
Proper Inclusion	\subset	$<$

When restricted to units within a given partition, \subseteq and \leq coincide, and so do \cap and $*$. We can think of \subseteq as the result of restricting \leq to the *natural units* picked out by the partition in question. We can think of set theory as amounting to the overlooking of the idea that there is a distinction between natural units and arbitrary unions. Set theory, indeed, derives all its power from this overlooking.

3.6. *Trees*

Philosophers since Aristotle have recognized that the results of our sorting and classifying activities can be represented as those sorts of branching structures which mathematicians nowadays call *trees*. For an example of a tree, see Figure 3, which represents the top level of the Foundational Model of Anatomy (FMA) ontology (see *FMA*, also Rosse and Mejino, 2007).

Trees are directed graphs without cycles (i.e., if we move along its edges, then we will always move down the tree and in such a way that we will never return to the point from which we started). They consist of nodes or vertices, and of directed edges, that connect the nodes. That the edges are directed means that the vertices connected by an edge are related to each other in a way that is analogous to an ordered pair; the relation between nodes is asymmetric or unidirectional. Here, we are interested specifically in rooted trees, that is, trees with a single topmost node to which all other vertices are connected, either directly or indirectly, via edges. In a rooted tree, every pair of vertices is connected by one and only one chain (or sequence of edges). We shall think of the directedness of an edge as proceeding down the tree from top to bottom (from ancestors to descendants). The connection between partitions and trees will now be obvious, as it is a simple matter to show that every finite partition can be represented as a rooted tree of finite depths, and vice versa (Mark, 1978).

Figure 3: *Highest-level Branches of the Tree Representing the FMA*

4. Granular Partitions in their Projective Relation to Reality

4.1. Projection

Partitions are more than just systems of units. They are constructed to serve as inventories or pictures or maps of specific portions of reality and, in this way, they are analogous to *pictures* or *windows* (Smith, 2001a). They are also analogous to propositions (*Elementarsätze*) as described by Wittgenstein in the *Tractatus* (1961). A proposition, for Wittgenstein, is built out of simple signs (names) arranged in a certain order. Each name, Wittgenstein tells us, stands in a projective relation to a corresponding object in the world; it cannot fail to strike its target. If a proposition is true, then its simple signs stand to each other within the proposition as the corresponding objects stand to each other in the world. It is in this sense that a true atomic proposition is a *picture* of a state of affairs in reality, as Wittgenstein puts it. That a proposition is a complex of names arranged in

a certain order is, in our present context, equivalent to the thesis that a partition is a complex of units arranged in a certain order.

A partition is a complex of units in its projective relation to the world (compare *Tractatus*, 3.12). This relation may be effected either directly by the user of the partition – for example, in looking through the units of the grid and recording what objects are detected on the other side – or indirectly, with the help of proper names or other referring devices such as systems of coordinates or taxonomic labels.

From the perspective of granular partition theory projection may fail, and a partition may be such that there are no objects for its units to project onto (like the partition cataloguing Aztec gods). Here, however, we are interested primarily in partitions which do not project out into thin air in this way. In what follows, we shall assume that a unique projection is defined for each partition. In a more general theory, we can weaken this assumption by allowing projections to vary with time while the partition remains fixed, for example (this is allowed in Smith and Brogaard, 2002). Such variation of projection for a fixed partition is involved in all temporally extended sampling activity. Consider, for example, what happens when we use a territorial grid of units to map the presence of one or more birds of given species in given areas from one moment to the next.

4.2. *Location*

If projection is successful, then we shall say that the object upon which a unit is projected is located in that unit. The use of the term ‘location’ reflects the fact that one important inspiration of our work is the study of location relations in spatial contexts. One motivating example of a location relation of the sort here at issue is the relation between a spatial object such as a factory building and the corresponding rectangular icon on a map. Other motivating examples are of a non-spatial sort and include the relation between an instance (Tibbles) and its kind (cat).

In what follows, we make the simplifying assumption that objects are located precisely at their units. That is, we will assume that the boundaries of the real-world objects correspond to the boundaries of the partitions through which we apprehend them. Compare the way in which your brother Norse is exactly located at the unit ‘Norse’ in your partition (list) of your family members. In a more general theory, we liberalize the location relation in such a way as to allow for partial or rough location as well (Casati and Varzi, 1995; Bittner and Stell, 1998); as, for example, between

a factory building and the corresponding square formed by the grid on a map.

4.3. *Transparency*

When projection succeeds, location is what results. Projection and location thus correspond to the two directions of fit – from mind to world and from world to mind – between an assertion and the corresponding truthmaking portion of reality. (For seminal work in this area, see Searle, 1983; compare also Smith, 1999.) Projection is like the relation which holds between your shopping list and the items which, if your shopping trip is successful, you will actually buy. Location is like the relation which obtains between the items you have bought and the new list your mother makes after your return, as she checks off those items which you have in fact succeeded in bringing back with you.

The formula ‘ $L(o, z)$ ’ abbreviates: object o is located at unit z . (And again: where this is required we can write ‘ $L_A(o, z)$ ’ for: o is located at z in partition A .) Location presupposes projection; an object is never located in a unit unless that object has already been picked out as the target of the projection relation associated with the relevant partition. But *successful* projection – by which is meant the obtaining of the projection relation between a unit and an object – also presupposes location, so that where both location and projection obtain they are simply the converse relations of each other.

We have now reached the point where we can formulate the first of our master conditions on partitions from the perspective of theory (B):

MB1: If object o is located at unit z , then unit z projects onto object o .

MB2: If unit z projects onto object o , then object o is located at unit z .

(Successful) projection and (successful) location are simple converses of each other. MB1 and MB2 tell us that a partition projects a given unit onto a given object if and only if that object is located in the corresponding unit. Very many partitions – from automobile component catalogues to our maps of states and nations – have this quality without further ado.

We shall call partitions which satisfy MB1 and MB2 *transparent partitions*, a notion which we can define as follows:

DTr: Partition A is *transparent* if and only if, for every unit z of partition A and for every object o:

- (i) if z projects onto o, then o is located at z, and
- (ii) if o is located at z, then z projects onto o.

MB1 and MB2 jointly ensure that objects are actually located at the units that project onto them. Notice that, according to our definition, a transparent partition may still have empty units. Such units may be needed to leave room for what may be discovered in the future or to cover up for temporary lapses in memory.

4.4. *Functionality Constraints (Constraints Pertaining to Correspondence to Objects)*


4.4.1. *Projection is Functional: The Confused Schoolboy*

The property of transparency is still rather weak. Thus, transparency is consistent with ambiguity on the side of the units in relation to the objects they target. Such is the case where one unit projects onto two distinct objects. An example of the sort of problem we have in mind is the partition created by a lazy schoolboy studying the history of the Civil War in England. This partition has just one unit labeled ‘Cromwell’, and so it does not distinguish between Oliver and his son Richard.

Although such ambiguous units do sometimes exist, in an ideal scenario they should be rectified when they are discovered. To eliminate such ambiguity, we lay down a requirement that each partition must be such that its associated projection is a *functional* relation:

MB3: If unit z projects onto object o_1 and onto object o_2 , then o_1 and o_2 are identical.

Figure 4: *Transparent Partitions in which Projection is not Functional (left) and Location is not Functional (right)*



For partitions satisfying MB3, units are projected onto single objects. Consider the left part of Figure 4. The dotted arrow can occur in partitions satisfying merely MB1 and MB2 but not in partitions also satisfying MB3. Notice, though, that projection might still be a partial function, since MB3 does not rule out the case where there are empty units.

4.4.2. *Location is Functional: The Morning Star and the Evening Star*

Consider a partition with a maximal unit labeled ‘heavenly bodies’ and three subunits labeled ‘The Morning Star’, ‘The Evening Star’, and ‘Venus’, respectively. As we know, all three subunits project onto the same object. This partition is perfectly consistent with the conditions we have laid out thus far. Its distinct subunits truly, though accidentally, project onto the same object. However, a good partition should clearly be one in which such errors of duplication of representational units are avoided.

Partitions manifesting the desired degree of correspondence to objects in this respect must in other words be ones in which location, too, is a *functional* relation:

MB4: If the same object is located in unit z_1 and in unit z_2 , then units z_1 z_2 are identical.

MB4 ensures that location is a function, i.e., that objects are located at single units (one rather than two). Consider the right part of Figure 4. The dotted arrow can occur in partitions satisfying MB1 and MB2, not however in partitions also satisfying MB4. As MB3 rules out co-location, so MB4 rules out co-projection. Note that natural analogues of co-location and co-projection cannot be formalized within a set-theoretic framework.

5. *Correspondence of Mereological Structure*

Even in the presence of MB3 and MB4, MB1 and MB2 tell us only that if a unit in a partition projects upon some object, then that object is indeed located in the corresponding unit. They do not tell us what happens in case a unit fails to project onto anything at all. Thus, MB1–4 represent only a first step along the way towards an account of correspondence to reality for partitions. Such correspondence will involve the two further dimensions of *structural mapping* and of *completeness*.

5.1. *Recognizing Mereological Structure*

An object o is *recognized* by a partition if and only if the latter has a unit in which that object is located (Smith and Brogaard, 2000). Intuitively, recognition is the partition-theoretic analogue of the standard set-membership relation. To impose a partition on a given domain of reality is to *foreground* certain objects and features in that domain and trace over others. Partitions are granular in virtue precisely of the fact that a partition can recognize an object without recognizing all its parts.


Partitions are designed to reflect the part-whole structure of reality through the fact that the units in a partition are themselves such as to stand in relations of part to whole. Given the master conditions expressed within the framework of theory (A) above, partitions have at least the potential to reflect the mereological structure of the domain onto which they are projected. In felicitous cases this potential is realized.

That we distinguish between the recognition (foregrounding, selection) of objects, on the one hand, and the reflection of mereological structure, on the other hand, is not an arbitrary matter. In Tractarian semantics, we distinguish between projection and isomorphism. The former obtains where we have some representation (for example a list or a map) which is intended to capture the entities in some domain of reality, while the latter obtains where this intention is fulfilled, so that there is a one-one correspondence between the units of the representation and the entities in the represented domain. In set theory we distinguish, for any given set, between a domain of elements and the set-theoretic structure imposed on this domain. Just as it is possible to have sets consisting entirely of *Urelemente* (together with a minimal amount of set-theoretic packaging), so it is possible to have partitions built exclusively out of minimal units (and one maximal unit). Such partitions amount, simply, to lists of the things that are recognized by their units, with no mereological structure on the side of these objects being brought into account.

Figure 5 (a) and 5 (b) represent partitions consisting of two minimal units z_1 and z_2 projecting onto objects o_1 and o_2 . Case (a), a simple list, is unproblematic. Case (b) we shall also allow as unproblematic. This is in keeping with the notion that minimal units are the (relative) atoms of our system, and we take this to mean that they should be neutral with regard to any mereological structure on the side of their objects. An example of type (b) would be a list of regions represented at a conference to discuss measures against terrorism, a conference including representatives from

both Germany and Bavaria. Here we are not concerned about the fact that Bavaria is a part of Germany.

Figure 5: *Transparent Partitions with More or Less Desirable Properties*



Cases like (c), in contrast, represent projections in which, intuitively, something has gone wrong. All three cases satisfy the master conditions we have laid down thus far, for the latter allow both for disjoint units to be projected onto what is not disjoint (b) and also for disjoint objects to be located in units which are not disjoint (c). Cases like (c) seem to fly in the face of a fundamental principle underlying the practice of hierarchical classification, namely, that objects recognized by species lower down in a hierarchical tree should be included as parts in whatever is recognized by the genera further up the tree. To exclude cases like (c), we shall impose a condition that mereological structure within a partition should *not misrepresent the mereological relationships* between the objects onto which the corresponding units are projected. We first of all define the following relation of *representation of mereological structure* between pairs of units:

DS1: Units z_1 and z_2 represent the mereological structure of the objects onto which they project if and only if, for objects o_1 , and o_2 : if o_1 is located in z_1 and o_2 is located in z_2 , and if z_1 is a subunit of z_2 , then o_1 is part of o_2 .

If z_1 is a subunit of z_2 then any object recognized by z_1 must be a part of any object recognized by z_2 .

DS2: A partition is mereologically structure-preserving if and only if: each pair of units within the partition satisfies DS1.


We can now impose a new master condition:

MB5: All partitions are *mereologically structure-preserving* in the sense of DS2.

Note that even MB5 is still very weak. Its effect is entirely negative, since it merely ensures that partitions do not misrepresent the mereological relationships between their objects. Partitions might still be entirely blind to (i.e. they may trace over) such relationships. Two minimal units might project onto objects which stand to each other in any one of the possible mereological relations (identity, proper parthood, disjointedness, overlap), and all pairs of units are likewise neutral as to the mereological relations between the objects onto which they are projected, provided only that they do not stand to each other in the subunit relation. This means that, given such units, we are entitled to infer nothing at all about the mereological relations among the corresponding objects.

Consider, for example, a partition that contains two units that recognize, respectively, mammals and whales. Suppose that this is a partition constructed at a time when the status of whales as mammals was not yet recognized. The unit labeled *whales* is not, then, included as a subunit of the unit labeled *mammals*. But the partition can still satisfy our conditions laid down so far. This is so, for example, if the unit that recognizes whales is a subunit of the unit recognizing animals but not a subunit of any other subunit of the unit recognizing animals (Partition A_1 in Figure 6). If the unit that recognizes whales were also a subunit of the unit that recognizes fish, for example, then the partition would misrepresent the mereological relationship between these two species and so violate MB5 (Partition A_2 in Figure 6).

Figure 6: Partition A_1 does not misrepresent the mereological structure of the underlying domain. Partition A_2 places whales incorrectly in relation to fish and mammals.



Partitions may trace over mereological relationships between the objects they recognize, but MB5 is strong enough to ensure that, if a partition tells us something about the mereological relationships on the side of the

objects which it recognizes, then what it tells us is true. Notice that partition A_2 still satisfies MB1–4.

Consider a domain of objects consisting of two regions, x and y , that overlap in the region v , so that the product of x and y is v , with v being a proper part of x and of y . (In symbols: $x * y = v$ with $v < x$ and $v < y$.) Consider now a partition with units z_1 and z_2 recognizing x and y , respectively, so that x is located in the first and y in the second unit: $L(x, z_1)$ and $L(y, z_2)$. Assume further that z_1 and z_2 do not stand in any subunit relations to each other. Only four possibilities regarding the representation of v now remain: (1) our partition does not recognize v at all; (2) it recognizes v but traces over its mereological relationships to x and y ; (3) it recognizes v through a subunit of z_1 but it traces over the fact that v could equally well be recognized by a subunit of z_2 ; (4) it recognizes v through a subunit of z_2 , but it traces over the fact that v could equally well be recognized by a subunit of z_1 . The fifth possibility – of allowing sub-units of both z_1 and z_2 to recognize v – is excluded by the tree structure of granular partitions.

Let x and y be two neighboring countries which disagree about the exact location of their common boundary and let v be the disputed area. The inhabitants of country x consider v to be part of x , the inhabitants of country y consider v to be part of y . Possibility (1) then corresponds to the view of some third country at the other side of the globe which recognizes the countries x and y but does not care about their border dispute. (2) corresponds to the view of an observer who recognizes that there is a disputed area but who is neutral about the status of the disputed area. (3) corresponds to the view of country x and (4) to that of country y .

Another example of case (2) is provided by Germany and Luxemburg, which overlap at their common border on the River Our. The river is part of both countries. Mapmakers normally have no facility to represent cases such as this, and so they either adopt the policy of not representing such common regions at all (the border is represented as a line which we cannot properly in this case imagine as being without thickness), or they recognize the region constituted by the river on the map but trace over its mereological properties. Larger-scale maps often embrace a third alternative, which is to *misrepresent* the relations between Germany and Luxenburg by drawing the boundary between the two countries as running down the center of the river.

5.2. *The Domain of a Partition*

That upon which a partition is projected is a certain domain of objects in reality (the term ‘domain’ being understood in the mereological sense, i.e. as a mereological sum of entities in some region of reality demarcated in some way, for example on the basis of our focus of interest). We shall conceive the domain of a partition as the mereological sum of the pertinent objects, those upon which the partition sets to work: thus it is stuff conceived as it is prior to any of the divisions or demarcations effected by the partition itself. The domains of partitions will comprehend not only individual objects and their constituents, but also groups or populations of individuals (for example biological species and genera) as well as their constituent parts or members. Domains can comprehend also extended regions (bounded continua) of various types. Spatial partitions, for example maps of land use or soil type (Frank *et al.*, 1997), are an important family of partitions with domains of this sort.

We are now able to specify what we mean by ‘domain of a partition’. Our representation of partitions as trees and our condition on reflection of structure (MB5) ensure that all partitions trivially reflect the fact that the objects recognized by their units are parts of some mereological sum. For MB5 is already strong enough to ensure that everything that is located at some unit of a partition is part of what is located at the corresponding maximal unit.

We can thus define the domain of a partition simply as the object (mereological whole) onto which its maximal unit is projected. By functionality of projection and location there can be only one such object.

DD: The domain of partition A is identical to the object upon which its maximal unit is projected.

We now demand as a further master condition that every partition have a non-empty domain in the sense of DD:

MB6: For every partition A, there is an object x which is identical to the domain of partition A.

We then say that a partition *represents its domain correctly* if and only if MA1-5 and MB1-6 hold. Correct representations, as we see, can be highly partial.

5.3. *The Granularity of Granular Partitions*

A correct representation is not necessarily a complete representation. Indeed, since partitions are cognitive devices, and cognition is not omniscient, it follows that no partition is such as to recognize all objects. This feature of partitions is captured by the following condition:

FP: There is no partition which projects onto everything in the sense that its maximal unit projects onto the universe as a whole and it has a subunit recognizing every object there is.

This condition does justice to the fact that the complexity of the universe is much greater than the complexity of any single cognitive artifact. This feature of partiality is captured already by our terminology of *granular* partitions. Partitions characteristically do not recognize the proper parts of the whole objects which they recognize.

It is the units of a partition which carry this feature of granularity with them. Like singletons in set theory, they recognize only single whole units, the counterparts of set-theoretic elements or members. If a partition recognizes not only wholes but also one or more parts of such wholes, then this is because there are additional units in the partition which do this recognizing job. Consider, for example, a partition that recognizes human beings and has units that project onto John, Mary, and so forth. This partition does not recognize parts of human beings – such as John’s arm or Mary’s shoulder – unless we add extra units for this purpose. Even if a partition recognizes both wholes and also some of their parts, it is not necessarily the case that it also reflects the mereological relationships between the two.

Partitions are cognitive devices which have the built-in capability to recognize objects and to reflect certain features of those objects’ mereological structure and to ignore, or trace over, other features of this structure. We can now see that they can perform this task of tracing over in two ways, namely, (1) by tracing over mereological relations between the objects which they recognize, and (2) by tracing over, or failing to recognize, parts of those objects. Unless there is some smallest unit, (2) is a variety of tracing over that must be manifested by every partition. A third type of tracing over becomes apparent when we remember that partitions are partial in their focus, and thereby each partition traces over everything that lies outside its domain (here, we leave to one side the Spinoza

partition, namely, the *monad*, which consists of just one all-encompassing universal unit).

Consider a simple biological partition of the animal kingdom including just one single unit, projecting on the species dog (*canis familiaris*). Our definition of the domain of a partition and our constraint on functionality of projection imply that, besides the species dog, also your dog Fido, and also Fido's DNA-molecules, are parts of the domain of this partition. *But the latter are of course not recognized by the partition itself.*

Partition theory hereby allows us to define a new, restricted notion of parthood that takes granularity into account. This restricted parthood relation is an analogue of partition-theoretic inclusion, but on the side of objects:

DRP: x is a part of y relative to partition A if and only if x is recognized by a subunit of a unit in A which recognizes y .

From DRP we can infer by MB5 that x is a part of y also in the unrestricted or absolute sense.

The usual common-sense (i.e., non-scientific) partition of the animal kingdom contains units recognizing dogs and mammals, but no units recognizing DNA molecules. Relative to this common-sense partition, DNA molecules are not parts of the animal kingdom in the sense defined by DRP, though they are of course parts of the animal kingdom in the usual, mereological sense of 'part'.

6. *Structural Properties of Correct Representations*

What are some of the more fundamental varieties of those partitions which satisfy the master conditions set forth above? We can classify such partitions according to: (1) the degree to which they match the structure of the objects which they represent (i.e., their *structural fit*); (2) their degree of completeness and exhaustiveness with respect to their domain; (3) their degree of redundancy (the smaller the redundancy the more adequate the representation). (For more on this, see Ceusters and Smith, 2006; and Smith, Kusnierczyk, Schober, and Ceusters, 2006).

6.1. *Structural Fit*

We require of partitions that they at least not misrepresent the mereological structure of the domain they recognize. This leaves room for the possibility that a partition is merely neutral about (traces over) some or all aspects of the mereological structure of its target domain. Taking this into account, we can order partitions according to the degree to which they actually do represent the mereological structure on the side of the objects onto which they are projected. At the maximum degree of structural fit, we have those partitions which completely reflect the mereological relations holding between the objects which they recognize.

Such a partition satisfies a condition which is the weak converse of MB5:

CM: If object o_1 is part of object o_2 , and if both o_1 and o_2 are recognized by a partition A, then the unit at which o_1 is located is a subunit of the unit at which o_2 is located.

A partition satisfying CM is *mereologically monotonic*: it recognizes all the restricted parthood relations obtaining in the pertinent domain of objects. A very simple example is given by a flat list projected one-for-one upon a collection of disjoint objects.

6.2. *Completeness*

So far we have allowed partitions to contain empty units, i.e., units that do not project onto any object. We now consider partitions which satisfy the constraint that every unit recognizes some object:

CC: If z is a subunit of partition A, then there is some object o which is located at z .

We say that partitions that satisfy CC *project completely*. Of particular interest, however, are partitions that project completely and in such a way that projection is a total function (partitions which satisfy both MB3 and CC). An example is a map of the United States representing its constituent states (with a whole for the District of Columbia). There are no non-states within the territory projected by such a map and every unit projects uniquely onto just one state.

6.3. Exhaustiveness

So far we have accepted that there may be objects in our target domain which are not located at any unit. This feature of partitions is sometimes not acceptable: governments want *all* their employed citizens to be located in some unit of their partition of taxable individuals. They want their partitions to satisfy a completeness constraint to the effect that every object in the domain is indeed recognized. In this case we say that location is *complete*, or that the partition *exhausts* its domain. We might be tempted to apply the following axiom to capture the exhaustiveness constraint:

(*) If object o is included in the domain onto which partition A projects, then there is a unit z at which o is located.

However, this condition is unrealizable because not every object is recognized in every partition: the tax authorities do not (as of this writing) want to tax the separate molecules of their citizens, and so the partitions of reality which they employ do not recognize these molecules.

It will in fact be necessary to formulate several restricted forms of exhaustiveness, each one of which will approximate in different ways the condition expressed in (*).

One such exhaustiveness condition might utilize a sortal predicate (schema) φ that singles out the kinds of objects our partition is supposed to recognize (for example, in the case of the just-mentioned partition, the predicate *being a taxable individual*). We now demand that the partition recognize all of those objects in its domain which satisfy φ :

CE_{φ} If object o is included in the domain onto which partition A projects, and if o satisfies requirement φ , then there is a unit z at which o is located.

We can very simply use any predicate to define a partition over any domain, by the following definition:

Object o is located in unit z of partition A if and only if:

- (i) o is properly included in the domain onto which A projects and
- (ii) o satisfies requirement φ .

Hence CE_{φ} entails the completeness of one partition *relative to* another.

The intuition behind constraints like CE_{φ} is to catch everything above a certain resolution at least once. If you want your partition to rule out


locating an item in two separate units, so that no item is counted twice, then you want the partition to satisfy CE_ϕ and MB4. You may also want to make sure that your partition does not include empty units, which you rule out by imposing condition CC. For partitions that satisfy the criteria CE_ϕ , and MB1-5, and CC projection and location are total functions (relative to the selected predicate ϕ) and one is the inverse of the other. Notice that neither of the following holds:

(**) if MB4 and CE_ϕ and CC, then MB3

(***) if MB3 and CE_ϕ and CC, then MB4

Counterexamples are given in Figure 7 (a) and (b), respectively, where each depicted object is assumed to satisfy ϕ .

Figure 7: *Functionality of Projection and Location are Independent of Completeness and Exhaustiveness.*



6.4. Comprehension Axioms

The following is the partition-theoretic equivalent of the unrestricted set-theoretic comprehension axiom. For each predicate ϕ there is a partition $A(\phi)$ whose location relation is defined as follows:

There is a unit in partition A in which object o is located if and only if o satisfies ϕ .

Under what conditions on ϕ can this be allowed? One type of restriction that is relevant to our purposes would allow ϕ to be unrestricted but would affirm additional restrictions on objects, for example in terms of spatial location. Thus we might define a family of spatial partitions $A(\phi, r)$, where r is some pre-designated spatial region, in such a way that:

There is a unit in partition A in which object o is located if and only if:

- (i) o satisfies ϕ , and
- (ii) o is spatially located in region r .

Something like this is in fact at work in the partitions used by epidemiologists, ornithologists, and others who are interested in (types of) objects at specific sites.

6.5. *Redundancy*

Partitions are natural cognitive devices used by human beings to serve various practical purposes. This means that partitions normally will be called upon to avoid certain sorts of redundancy. Here, we distinguish what we shall call *correspondence redundancy* and *structural redundancy*. Correspondence redundancy occurs where we have too many disjoint units in a partition. Structural redundancy occurs where two or more non-disjoint units are present which project onto the same portion of reality. Necessarily, empty units (such as would be defined by the predicate ‘entity that is not identical with itself’) represent one type of correspondence redundancy which is excluded by condition CC. Another type of correspondence redundancy is present in a partition with two distinct units whose labels would tell us *ex ante* that they must necessarily project upon the very same object. Clearly, and most simply, a partition should not contain two distinct units with identical labels.

The following case is not quite so trivial. Consider a partition with a unit labeled *vertebrates* which occurs as a subunit of the unit labeled *chordates* in our standard biological classification of the animal kingdom. Almost all chordates are in fact vertebrates. Suppose (for the sake of the illustration) that biologists were to discover that all chordates must be vertebrates. Then such a discovery would imply that, in order to avoid structural redundancy, they would need to collapse into one unit the two units (of chordates and vertebrates) which at present occupy distinct levels within their zoological partitions. A constraint designed to rule out such structural redundancy would be:

CR: A unit in a partition never has exactly one immediate descendant.

This rules out partition-theoretic analogues of the set theorist’s $\{\{a\}\}$.

7. *Fullness and Cumulativity*

Thus far we have distinguished completeness and exhaustiveness. We now introduce a third type of comprehensiveness factor for partitions, which is needed for ensuring that the successive levels of the partition relate to each

other in the most desirable way. We can initially divide this third type of completeness into two sub-types, namely, *fullness* and *cumulativeness*. Fullness requires that each unit z have enough subunits (which are immediate descendants) to fill out z itself. Cumulativeness requires that these immediate subunits be such that the objects onto which they are projected are sufficient to exhaust the domain onto which the containing unit is projected. Fullness, accordingly, pertains to theory (A), cumulativeness to theory (B).

Non-fullness and non-cumulativeness represent two kinds of shortfall in the *knowledge* embodied in a partition. Non-fullness is the shortfall which arises when a unit has insufficiently many subunits within a given partition (for instance it has a unit labeled *mammal*, but no subunits corresponding to many of the species of this genus). Non-cumulativeness is the shortfall which arises when our projection relation locates insufficiently many objects in the units of our partition, for example when I strive to make a list of the people that I met at the party yesterday, but leave out all the Welshmen. Fullness and cumulativeness are satisfied primarily by artificial partitions of the sorts constructed in database environments. In the remainder of this section we assume, for the sake of simplicity, that there are no redundancies in the sense of CR.

7.1. *Fullness*

Consider a partition with three units labeled: *mammals*, *horses*, and *sheep*. This partition is transparent, by our definition (DTr); but falls short of a certain sort of ideal completeness. We can express the problem as follows.

In set theory, if a collection of subsets of some given set forms a partition of this set in the standard mathematical sense, then these subsets are (1) mutually exhaustive and (2) pairwise disjoint (the latter meaning that the subsets have no elements in common). An analogue of condition (2) holds for minimal units in our present framework, since minimal units are always mereologically disjoint (they cannot, by definition, have subunits in common). Condition (1) however does not necessarily hold within the framework of partition theory. This is because, even where the partition-theoretic sum of minimal units is identical to the maximal unit, the minimal units still do not necessarily exhaust the partition as a whole. The mereological sum (+) of units is, we will recall, in general smaller than their partition-theoretic sum (\cup).

Thus, we can define a unit as *full* within a given partition if its subunits are such that their mereological sum and their partition-theoretic sum coincide.

DFull: Unit z is *full* if and only if: the mereological sum of its subunits is identical to the partition-theoretic sum of its subunits.

However, DFull does not suffice to capture the intended notion of fullness for partitions. To see the problem, consider the partition consisting of

<i>cells</i>			
<i>prokaryotic</i>		<i>eukaryotic</i>	
<i>nucleoid</i>	<i>ribosomes as only organelles</i>	<i>nucleus</i>	<i>membrane bound organelles</i>

The unit in the top row satisfies DFull, but it is not full relative to all of its subunits, since the mereological sum of the units *nucleoid* and *ribosomes as only organelle* is not identical to the unit *prokaryotic*. The problem arises because, if x is mereologically included in y , then the mereological sum of x and y is y ; and if x is partition-theoretically included in y , then the partition-theoretic sum of x and y is y . From this, it follows that only the immediate subunits of a given unit z_1 contribute to its mereological and partition-theoretic sums.

This, however, tells us what we need to take into account in defining what it is for a unit to be full relative to all its subunits within a given partition A , namely that each of its constituent units must be full relative to its immediate descendents. This yields:

DFull*: Unit z_1 is a full* unit of partition A if and only if, for every z : if z is a partition-theoretic part of z_1 , then z is either a full unit of A or a minimal unit of A .

Here minimal units have been handled separately because they do not have subunits. One can see that, while *cells* in the mentioned partition are full, they are not full*, because the units *prokaryotic* and *eukaryotic* are neither full nor minimal.

Thus far we have defined fullness for units. We can now define what it means for a *partition* to be full, as follows:

DFull2: A partition is full if and only if all its non-minimal units are full (or, equivalently, all its units are full*).

Notice that full partitions might in principle contain empty units, which may or may not have subunits.

7.2. *Cumulativeness*

Cumulativeness plays the same role in theory (B) which fullness plays in theory (A). The intuitive idea is as follows: a unit is cumulative relative to its immediate subunits if the mereological sum of the *objects onto which those subunits project* (for short: the units' *projection*) is identical to the projection of their partition-theoretic sum. For non-empty and non-minimal units with at least two immediate subunits we define:

DCu1: A unit z is *cumululative* if and only if: the partition-theoretic sum of its projection is identical to the mereological sum of its projection.

One can see that, under the given conditions, it is the case that the projection of a unit's partition-theoretic sum is identical to the projection of the unit. Consequently:

If a unit z is cumulative, then the mereological sum of its projections is identical to z , and vice-versa.

Again, the cumulative condition ensures that z_1 is cumulative relative to its immediate subunits. In order to ensure cumulativeness of a unit with respect to all its subunits, we define:

DCu2: A partition is cumulative if and only if all its units are cumulative.

DCu*: A partition is cumulative* if and only if all its units are either cumulative or minimal.

7.3. *Equivalence of Fullness and Cumulativeness*

From the definitions above, it follows that cumulative partitions for which CC and MB1-5 hold are full and that they do not contain empty units. To see why this is so, consider a non-minimal unit z_1 of which it holds that its partition-theoretic sum is identical to its projection.

We need to show that in this case z_1 itself is identical to its partition-theoretic sum. Since by MB5 projection does not distort mereological structure, two cases need to be considered:

- (a) z_1 and its partition-theoretic sum are mereologically disjoint, which means that they trace over the mereological relationships between z_1 's projection and the projection of its partition-theoretic sum; and
- (b) z_1 's partition-theoretic sum is mereologically included in z_1 , and z_1 is mereologically included in z_1 's partition-theoretic sum, in which case the mereological relationships between z_1 's projection and the projection of its partition-theoretic sum are preserved.

In case (b), z_1 is identical with its partition-theoretic sum, as desired. Case (a) cannot occur, however, since projection is a one-one mapping; that is, distinct units project onto distinct objects (by CC and MB1-4) and, therefore, the two distinct and disjoint units z_1 and its partition-theoretic sum cannot project onto one and the same object, as is required by our assumption that the projection of z_1 is identical to the partition-theoretic sum of z_1 's projection.

We can also show that, in the opposite direction, full partitions for which CC and MB1-5 hold are cumulative. To see this, assume that z_1 is a non-minimal unit in a full partition A. We then need to show that if z_1 is identical to its partition-theoretic sum, then z_1 's projection is identical to the projection of its partition-theoretic sum. Assume that z_1 is identical to its partition-theoretic sum. We need to consider two cases:

- (a) z_1 's partition-theoretic sum is a unit in partition A;
- (b) z_1 's partition-theoretic sum is not a unit in partition A.

In case (a), the following holds: By MB5 we have: z_1 's projection and the projection of its partition-theoretic sum are mereologically included in each other. Since the underlying unit-structure is full, the projection of z_1 's partition-theoretic sum is identical to the partition-theoretic sum of z_1 's projection. Thus, z_1 's projection and the partition-theoretic sum of its projection are mereologically included in one another. Hence they are, as desired, identical.

In case (b), z_1 's partition-theoretic sum is not a unit in the partition A. Consequently it is as if there is some extra entity x (which we will characterize in the next section as what we shall call 'empty space') which, together with z_1 's partition-theoretic sum, sums up to z_1 . Consequently z_1 and its partition-theoretic sum cannot be identical. This, however, contradicts our assumption of fullness. Therefore case (b) cannot occur.

It follows that the notions of fullness and cumulateness are logically equivalent for completely projecting partitions. We can accordingly distinguish just two classes of such partitions:

(1) Full and cumulative

Sample partition: a list of the 50 US states, divided into two sub-lists: the contiguous 48, the non-contiguous 2.

Objects: the states themselves.

Projection: the obvious 'Utah'–Utah projection relation.

(2) Non-full and non-cumulative

Sample partition: you have a terrible hangover, and your accounting of the people at the party consists of four units: *John, Mary, the Irish, the Welsh*. As it happens, Sally is Scottish.

A cumulative partition A is also exhaustive in the sense of (CE_φ) , with φ requiring that there be a unit which is a part of A and which projects onto an object. Full partitions are also mereologically monotonic (CM). To see this, assume that it holds that z_1 's projection is mereologically included in z_2 's projection. Two cases need to be considered:


- (a) z_1 and z_2 are mereologically disjoint and trace over the relationships between their respective projections; or
- (b) z_1 and z_2 are not mereologically disjoint in which case, by MB5, the mereological relations are preserved.

Consider case (a). If z_1 and z_2 are not subunits of each other we have their mereological sum not being equivalent to z_2 . Therefore, by fullness, we have z_1 's and z_2 's partition-theoretic sums, when mereologically added, not equaling z_2 's partition-theoretic sum. Since fullness implies cumulateness, we have the partition-theoretic sum of z_1 's projection, when mereologically added to the partition-theoretic sum of z_2 's projection, not being equal to the partition-theoretic sum of z_2 's projection; and the mereological sum of z_1 's projection and z_2 's projection not being equal to z_2 's projection. Hence we arrive at a contradiction of our previous assumption: z_1 's projection is *not* mereologically included in z_2 's projection. Therefore case (a) cannot occur. In case (b) preservation of mereological structure immediately yields z_1 being partition-theoretically included in z_2 , which is the desired result.

7.4. *Empty Space*

When a unit falls short of fullness, it successfully projects onto some given domain, but its subunits do not succeed in projecting onto the entirety of this domain. It is then as if there is some extra but invisible component in the unit, in addition to its subunits. We shall call this additional component ‘empty space’ (noting that the term ‘empty’ here has a quite different meaning from what it has in the phrase ‘empty unit’ as used above). Consider the partition depicted in Figure 8:

Figure 8: *A Partition with Empty Space.*



Here, the empty space is that part of z_3 which is not occupied by z_1 and z_2 . Notice that this empty space is a component of z_3 but it is neither itself a unit nor is it made up of units. Empty space is that part of a unit which is not covered by its subunits; it is a zone within a unit where no subunits are to be found, reflecting, for example, the fact that many taxonomies are incomplete at lower levels.

Another characteristic of empty space is that there must be something that potentially fills it. In our case this means: more subunits. Empty space is *inert* in the sense that it does not project onto anything. Empty space is normally hidden to the user of the partition in which it exists, for otherwise this user would surely have constructed a fuller partition. In some cases however a user might deliberately accept empty space in order to have the means of acknowledging that something has been left out. Alternatively, the existence of empty space in a given partition might be brought to the attention of the user. We point in a certain direction and ask: What is *there*? The theory of empty space hereby serves as the starting-point for an ontology of questions (Schuhmann and Smith, 1987); empty space corresponds to a hole in our knowledge.

7.5. Fullness and Emptiness

Recall that, if z_1 is a full unit, then it is identical to its partition-theoretic sum, and vice versa. This means that, if z_1 is not a full unit, then its partition-theoretic sum is properly included in it; and that, if this is the case, then it is not a full unit. Consequently we can define what it means for x to be the empty space of the unit z_1 as follows. We first of all define x fills z_1 :

DFills: Empty space x fills unit z_1 if and only if:

- (i) z_1 is not a full unit, and
- (ii) z_1 is identical the mereological sum of x and z_1 's partition-theoretic sum.

The empty space in z_1 is then z_1 's smallest filler, and we define:

DES: x is the empty space in z_1 if and only if

- (i) x fills the space not occupied by the subunits of z_1 and
- (ii) x is disjoint from all subunits of z_1 .

We note, in passing, that minimal units, on the basis of the definitions above, are either empty or they are completely made up of empty space. ('Minimal' means: there is no further knowledge available, within a given partition, as concerns the objects onto which minimal units are projected.) Intuitively it is clear that DES determines the empty space of a unit uniquely. Consider x and y , which are both empty spaces in unit z as defined by DES. Each then fills the space not occupied by z 's subunits, and are disjoint from all of z 's subunits. Thus x and y are identical.

7.6. Empty Space and Knowledge

The presence or absence of empty space is a dimension of a granular partition that is skew to the dimension pertaining to the existence of empty units. An empty unit is a unit that fails to project. Empty space is that in a partition which leaves room for the addition of new knowledge.

Figure 8 depicts a partition of the animal kingdom consisting of three units, where z_3 recognizes the animal kingdom as a whole, z_1 recognizes dogs, and z_2 recognizes cats. z_3 is the partition-theoretic sum of z_1 and z_2 , but the sum of the projections of z_1 and z_2 is less than z_3 . New units can be

inserted into the partition if new species are discovered (e.g., the species indicated by o_3).

There is at least one other sort of knowledge shortfall which we shall need to consider in a complete theory of partitions. This arises when there are missing levels within a partition-theoretic hierarchy. A partition of the phylum *arthropoda* which mentions all of its species and genera but leaves out the units for its classes (e.g. *insecta*) is an example of this sort of incompleteness.

8. Identity of Granular Partitions

As a step towards a definition of identity for partitions, Smith and Brogaard (2003) propose a partial ordering relation between partitions, which they define as follows:

Partition A is mereologically included in partition B if and only if every unit of z is also a unit of B.

They then define an equivalence relation on partitions as follows:

Partition A is identical to partition B if and only if:

- (i) A is mereologically included in B, and
- (ii) B is mereologically included in A.

Now, however, we can see that a definition along these lines will work only for partitions which are full. What, then, of those partitions which are equivalent in the sense of DE but not full? What are the relationships between the presence of empty and redundant units and the question of the identity of partitions? And what is the bearing on the question of identity of the phenomenon of empty space? Can partitions that have empty or redundant units be identical? Can partitions which are not full be identical?

The question of whether or not partitions that have empty or redundant units are identical cannot be answered without a theory of labeling. If corresponding empty units in two distinct partitions are to be considered as identical, they need to have at least the same labels. We can only address this question informally here.

Consider the partition of the people in your building according to *number of days spent behind bars*. You can construct this partition prior to undertaking any actual inquiries as to who, among the people in your building, might be located in its various units. Thus, even before carrying

out such inquiries, you can know that this is a more refined partition than, for example, the partition of the same group of people according to *number of years spent behind bars*. The two partitions are distinct, and they will remain distinct even if it should turn out that none of the people in your building has spent any time at all in jail. In both cases, all the people in your building would then be located in a unit labelled *zero*, and all the other units in both partitions would be empty. Yet the two partitions would be nonetheless distinct, not least because their respective maximal units would have different labels.

We can now return, briefly, to our question whether partitions that are neither full nor cumulative can be said to be identical. One approach to providing an answer to this question would be to point out that, even though two partitions are outwardly identical, they might still be such that there are different ways to fill the corresponding empty space. Suppose we have what are outwardly the same biological taxonomies used by scientists in America and in Australia at some given time, both with the same arrays of empty units. Suppose these partitions are used in different ways on the two continents, so that, in the course of time, their respective empty space gets filled in different ways. Were they still the *same* taxonomy at the start?

9. Conclusions

This essay is a first step towards a formal theory of granular partitions. We defined master conditions that need to be satisfied by every partition. These master conditions fall into two groups, namely, (A) master conditions characterizing partitions as systems of units, and (B) master conditions describing partitions in their projective relation to reality.

At the level of theory (A) partitions are systems of units that are partially ordered by the subunit relation. Such systems of units can always be represented as trees; they have a unique maximal component and they do not have cycles in their graph-theoretic representations. But partitions are more than just systems of units. They are also cognitive devices that are directed towards reality.

Theory (B) takes this latter feature into account by characterizing partitions in terms of the relations of projection and location. Units in partitions are projected onto objects in reality. Objects are located at units when projection succeeds. To talk of granular partitions is to draw attention to the fact that partitions are in every case selective; even when they recognize some objects, they will always trace over others.

Partitions are also capable of reflecting the mereological structure of the objects they recognize through a corresponding mereological structure on the side of their array of units. This does not mean, however, that all partitions actually do reflect the mereological structure of the objects they recognize. For an important feature of partitions is that they are also capable of tracing over (ignoring) mereological structure.

Our discussion of granularity showed that partitions have three ways of tracing over mereological structure: (1) by tracing over mereological relations between the objects which they recognize; (2) by tracing over the parts of such objects; and (3) by tracing over the wholes which such objects form. The tracing over of parts is (unless mereological atomism is true) a feature manifested by every partition, for partitions are in every case *coarse grained*. The tracing over of wholes reflects the property of granular partitions of foregrounding selected objects of interest within the domain onto which they are projected and of leaving all other objects in the background where they fall in the domain of unconcern.

And the relevance to this book? Nearly everything that has been said about ontologies and their relation to reality in this volume can be illuminated by conceiving ontologies as granular partitions standing in a relation of projection to the entities their terms denote.

Chapter 7: Classifications

Ludger Jansen

It has long been a standard practice for the natural sciences to classify things. Thus, it is no wonder that, for two and a half millennia, philosophers have been reflecting on classifications, from Plato and Aristotle to contemporary philosophy of science. Some of the results of these reflections will be presented in this chapter. I will start by discussing a parody of a classification, namely: the purportedly ancient Chinese classification of animals described by Jorge Luis Borges. I will show that many of the mistakes that account for the comic features of this parody appear in real-life scientific databases as well. As examples of the latter, I will discuss the terminology database of the National Cancer Institute (NCI) of the United States, the *NCI Thesaurus*.

1. *Chinese Animals, or How to Make a Good Taxonomy*

In a certain Chinese Encyclopedia, the *Celestial Emporium of Benevolent Knowledge*, as Jorge Luis Borges tells us (1981), the following taxonomy of animals can be found:

- (1) those that belong to the Emperor
- (2) embalmed animals
- (3) trained animals
- (4) suckling pigs
- (5) mermaids
- (6) fabulous animals
- (7) stray dogs
- (8) those animals included in the present classification
- (9) animals that tremble as if they were mad
- (10) innumerable animals
- (11) animals drawn with a very fine camelhair brush
- (12) others
- (13) animals that have just broken a flower vase
- (14) animals that from a long way off look like flies

This taxonomy is a sophisticated piece of literature. It is also a good example of a bad taxonomy. For the sake of brevity, I will call Borges's taxonomy 'CAT' for 'Chinese Animal Taxonomy'. What lessons can we learn from CAT? Here are some of the rules for good and useful taxonomies, which CAT contravenes:

- **Ontological Grounding:** Good taxonomies classify things on the basis of traits belonging to those things. This precludes meta-types such as type (12): *others*. Things do not belong to the *other* group because they have some particular trait (of being other). Similarly, (14) does not classify things on the basis of traits belonging to those things themselves, but on the basis of their appearance to an observer.
- **Structure:** Good taxonomies take into account the fact that types of things have subtypes: for example, in biology there are genera and species. In CAT, however, all types have equal standing. It could be argued that mermaids are fabulous animals, in which case (5) would need to be rendered as a subtype of (6).
- **Disjointness:** If we have such a hierarchy of types and subtypes, then anything that instantiates a subtype also instantiates the type of which it is a subtype. For example, in biological systematics, every animal that is a horse is also a mammal. However, types on the same level of biological classification should be disjoint: no animal is both a mammal and a reptile, or both a vertebrate and an invertebrate. CAT's types, however, do not meet this criterion: Type (1) animals that belong to the Emperor probably include trained animals belonging under heading (3) as well.
- **Exhaustiveness:** Good taxonomies subsume all the entities they purport to subsume. At times this can be difficult to achieve, as in the biological sciences where new species are often discovered in the course of empirical research. CAT, however, seems to be far from exhaustive, if we ignore the fact that we can put any animal whatsoever under heading (12), *others*. Sometimes, exhaustiveness and disjointness are grouped together as the *jointly exhaustive and pairwise disjoint* (JEPD) criterion of classification.
- **No ambiguity:** Good taxonomies do not use terms ambiguously. Fabulous animals, pictures of animals, and dead animals, however, are not animals, at least not in the same sense that pigs or dogs are animals. For this reason, the headings (2), (5), (6), and (11) do not fit into this schema. What is more, painted animals are not animals, but rather paintings in which animals are represented.
- **Uniformity:** Good taxonomies have a well-defined domain. The traits by which they classify their objects should be of a uniform kind and be exemplified throughout the domain. CAT, however, draws on

the distinguishing traits of several different kinds at once. Heading (1) sorts animals according to their owners, (4) according to species membership, among other things, (7) according to species membership plus the lack of an owner, (9) according to behavior, (13) according to the effects of behavior, and (14) according to an animal's appearance to a remote observer.

- **Explicitness and precision:** Good taxonomies are explicit and precise. Headings such as (12), others, fulfill neither criterion.
- **No meta-types:** Good taxonomies avoid meta-types that come about through the classification process itself. In CAT, heading (8) is such a meta-type, and any animal belonging to CAT belongs under heading (8). If all animals belong to CAT, then all animals belong under (8). Thus every animal that belongs under (8) also belongs under headings (1)-(7) or (9)-(14). If an animal belongs to CAT but does not feature under these headings, this is no problem at all. It can also belong to CAT if it is a member of heading (8) alone. Heading (8) is a very peculiar heading for a taxonomy.

Classifications containing such types as (8) lead to problems that correspond structurally to the semantic paradox engendered by the sentence (T): 'This sentence is true'. (T) is indeterminate with regard to its truth value (that is, it is neither determinately true nor determinately false) because every truth value will fit. If we assume that it is true, what it says is the case, i.e. that it is true, and that is what is required for it to be a true sentence. But if we assume that it is false, then, as with any false sentence, what it says is not the case. Each of the two truth-values, true and false, can consistently be attributed to (T).

In the same manner, whether or not we classify animals that do not belong to other CAT-types under (8) can only be determined arbitrarily. A good classification system should not allow for this kind of arbitrariness concerning which objects fit under its types. Things get worse with CAT*, which we might call a Russellian version of CAT, containing (8*) 'Animals that do not belong to CAT*' instead of (8). A type like (8*) leads to problems that correspond, structurally, to Russell's antinomy or the liar paradox: if an animal belongs to types (1)-(8) or (9)-(14), then it belongs to CAT* and thereby does not belong to (8*). This is clear. But if an animal does not belong to these types, we encounter a paradoxical situation. For if an animal did not belong to (8*) either, it would not belong to any CAT*-type at all, and so would not belong to CAT*. Animals that do not belong

to CAT*, however, belong to (8*). If we suppose that the animal does not belong to the other types, it follows that, if something does not belong to (8*), then it belongs to (8*). But anything that belongs to (8*) belongs to CAT*. So the animal in question does *not* belong to (8*) after all. Classification systems should eschew such situations whenever possible.

2. *Medical Information Systems, or How to Make a Bad Taxonomy*

We have used CAT as a heuristic tool to point out some of the mistakes that can be made in the construction of a classification system. These mistakes appear, not only in literary parodies like CAT, but also in actual scientific practice. I will show this in the following, by discussing the *National Cancer Institute Thesaurus* (NCIT). This will provide the opportunity to discuss the abovementioned mistakes in greater depth, as well as to propose some ways of repairing them.

The National Cancer Institute in the United States created the NCIT to support its battle against cancer by developing an online controlled vocabulary for annotating and indexing information relevant to cancer research (Fragoso, *et al.*, 2004; see also Ceusters, Smith and Goldberg, 2005). It contains more than 110,000 expressions and 36,000 terms of importance to cancer research, including 10,000 types of medical findings and disorders, more than 5,000 anatomical kinds, upwards of 3,500 chemicals and medicines, and approximately 2,000 types of genes.

2.1. *Structuredness: Groups and Animals*

Whereas CAT is totally unstructured, the NCIT does have a hierarchy of supertypes and subtypes. Nevertheless, in many places the NCIT is unstructured, and it is sometimes structured incorrectly. Consider, for example, the NCIT entry ‘Subgroup’, which NCIT defines as a ‘subdivision of a larger group with members often exhibiting similar characteristics’. We should suppose that subgroups are groups, and this would indeed be implied by the NCIT definition of group, which is: ‘Any number of entities (members) considered as a unit’. But this link between ‘Subgroup’ and ‘Group’ – an important bit of structure – is missing from the NCIT.

This example, also, shows that the NCIT is sometimes structured incorrectly. For example, as the supertype of ‘Subgroup’ NCIT gives ‘Grouping’, which it defines as a ‘system for classifying things into groups

or the activity of putting things together in groups'. But, as philosophical tradition knows (see for example Aristotle, *Categories* 3), the definition of the supertype must also be applicable to all its subtypes. Thus from the definition of 'Grouping', and from the fact that *Group* is considered to be a subtype of *Grouping*, we get the following absurd conclusion, that a *subgroup* is either a *system* for classifying things into groups or an *activity* of putting things together in groups.

The NCI *Thesaurus*'s classification of animals is of similar quality to Borges's CAT. In the NCIT, the type *animal* splits into the subtypes *invertebrate*, *laboratory animal*, *vertebrate*, and *poikilotherms*. The subtypes *vertebrate/invertebrate* already present a problem, since they are an exhaustive division of all animals (and a division frowned upon by some biologists). Second, the artificial type *laboratory animal* stands out inappropriately when listed alongside the three natural classes, since laboratory animals do not comprise a natural kind. The subdivision appeals to traits of a range of different sorts. Finally, in reality *poikilotherms* is a subtype of *vertebrate* and, so, should not be classified at the same level as its supertype.

2.2. Disjunctiveness and Exhaustiveness: Patients

NCIT often contains subtypes which are not disjoint under the same supertype. An example is the entry *patient*. This entry has two subtypes: *cancer patient* and *outpatient*. These two entries are not disjoint, for many cancer patients are treated as outpatients. And naturally, these two subtypes are not an exhaustive classification of patients. There are many patients who are neither cancer patients nor outpatients. Normally, we would regard this example as a typical case of cross-classification, as there are two traits that an object could have independently of one another. Combined, these traits yield four classes of patients, as is presented in Figure 1:


Figure 1: *Four Classes of Patients: A Cross-classification*

PATIENTS	Outpatient? <i>Yes</i> .	Outpatient? <i>No</i> .
Cancer? <i>Yes</i> .	Outpatient with cancer	Inpatient with cancer
Cancer? <i>No</i> .	Outpatient without cancer	Inpatient without cancer

Classification systems are often constructed in such a way as to have the structure of an inverted tree, with a single highest-level root node and all

nodes beneath this root having at most one single parent node. This practice derives from the long tradition of the Porphyrian tree, named after the neo-Platonist Porphyry (ca. 234–304), whose introductory guide to Aristotle’s *Organon*, the *Isagoge*, presents the central headings of the classic Porphyrian tree as they appear in Figure 2. Such trees make it possible to construct definitions on the pattern of Aristotle: a species is defined according to its next highest type (the *genus proximum*), together with the specific traits which constitute the species (the *differentia specifica*). The stock example is still the definition of ‘human being’ as ‘rational animal’, citing both the proximate genus (‘animal’) and the specific difference that distinguishes human beings from animals of other kinds (‘rational’).

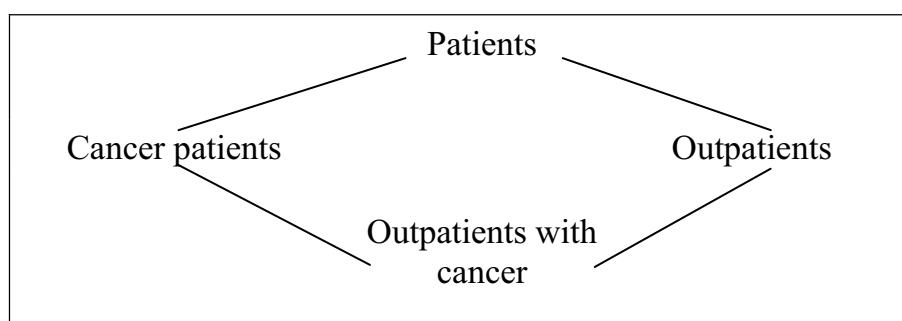
Figure 2: *The Structure of a Porphyrian Tree*



In information science, such tree structures are types of structured graphs. They flow in one direction, and the trees have a stem, the *genus ultimum*, from which increasingly finer branches split off, that finally end in the leaves or *species*. Taken together, all the ultimate kinds form the top-level ontology of an information system. In our stock example, the ultimate genus from which the species of human beings finally derives is normally assumed to be the category of substance or independent continuant (see chapter 8). Each element in such a tree (every node of the graph) has a unique supertype.

If we try to turn a cross-classification like the NCIT into a graph of this sort, then we face two problems. First, the uniqueness of a term's supertype is lost. Outpatients with cancer belong both to the supertype 'cancer patient' and to the supertype 'outpatient'. The branches of such a diagram no longer flow in a single direction. One element of the diagram can have multiple subtypes as well as multiple supertypes. Such situations are called *multiple inheritance* cases, since they allow us to produce diamond-formed structures like the example in Figure 3, in which the properties of the entities referred to by terms higher up in the hierarchy are inherited by the entities referred to by terms lower down along two or more distinct roots.

Figure 3: *An Example of Multiple Inheritance*



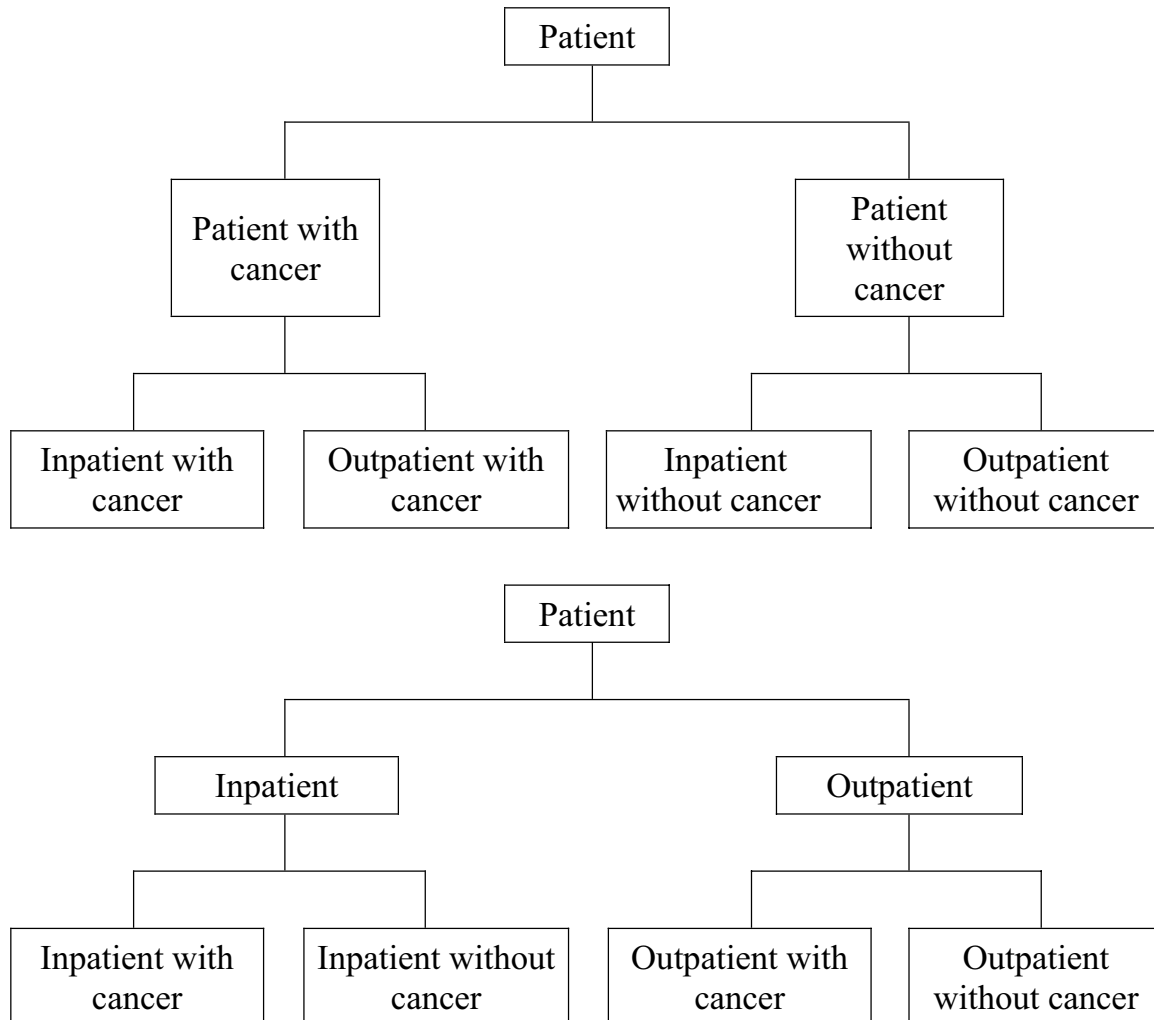
The second problem we face in such a situation is that, in order to construct a tree diagram after the fashion of Figure 4, we must determine which of these two traits should be considered prior in our classificatory hierarchy. In our classification, should we give priority the fact that the patient is an outpatient, or to the fact that he has cancer? To achieve a tree-structure, we must choose between the two.

Our choice between these two options would most likely be irrelevant to medical practice. But from the philosophical point of view, and from the point of view of ensuring consistency between different information systems (for example, in different medical specialties) such arbitrariness – and, thus, the possibility of making a random decision – is an unwelcome phenomenon, compounded by the fact that errors often result when distinct specification factors are combined within a single tree (Smith and Kumar, 2005). A cross-classification is based on the existence or nonexistence of two traits which are independent of one another. In the case of the patients in the NCIT, these are the questions: (1) for what is the patient being treated? (2) Is the patient staying overnight at the hospital?

The first question concerns the *reason* for the treatment, the second concerns an aspect of the *way in which he is treated*. Though both

questions are important for the doctor in the hospital, each answer comes from totally different categories (as we will see in Chapter 8 of this book), and should be strictly distinguished in a classification system.

Figure 4: *Two Alternative Tree Diagrams*



One possibility for separating these aspects of a patient from one another is to create a multi-dimensional (or multi-axis) classification system. This approach is used, for example, by SNOMED CT, the *Systematized Nomenclature of Human and Veterinary Medicine*, developed by the College of American Pathologists (see SNOMED). In its third version, SNOMED distinguishes eleven different axes (or traits by which to classify), which can be combined with 17 qualifications. Figure 5 lists some of SNOMED's semantic axes.

Not every disease representation requires each of these axes. But by appealing to multiple axes, an encephalitis virus in a forest ranger can be coded as: TX2000 M40000 E30000 J63230 where the part of the code

beginning with ‘T’ specifies the location of the disease, the part beginning with ‘M’ the body part affected, the part beginning with ‘E’ the cause of the disease (the virus), and the part beginning with ‘J’ the profession of the patient.

Figure 5: *Multi-dimensional Classification in SNOMED II*
(Dugas and Schmidt, 2003, 80)

Which morphological structure?	Morphology	M
Where is it situated?	Topography	T
What caused it?	Etiology	E
What is its effect?	Function	F
Which disease?	Disease	D
Which procedures have been applied?	Procedure	P
Connected with which profession?	Job	J

This correspondence of classificatory axes to kinds of questions is anticipated in the work of Aristotle, who uses terms for his categories which are taken mainly from interrogatory pronouns (Kahn, 1978, 227-278; cf. also next chapter).

2.3. *Uniformity: Laboratory Animals*

To classify patients according to both their cancer diagnosis and their status as outpatient leads to problems, not only with disjunctiveness and exhaustiveness, but it also violates the uniformity rule. Such a classification brings together distinguishing marks from different areas. This sort of violation is even more clearly manifest in the classification of laboratory animals in the NCIT. The importance of laboratory animals in cancer research is reflected in the variety of the twelve subtypes under the NCIT heading ‘Laboratory_Animals’. Some of these types reflect particular things that have happened to the animals in question. For example, according to the NCIT definition, a ‘Genetically_Engineered_Mouse’ is a ‘mouse that has been genetically modified by introducing new genetic characteristics to it’. Here, a DNA manipulation is given as the essence of a ‘Genetically_Engineered_Mouse’. Other types, like ‘Control_Animal’, reflect a certain role the animals take on within a certain experimental design:

Control_Animal NCI-GLOSS: the animals in a study that do not receive the treatment being tested. Comparing the health of control animals with the health of treated animals allows researchers to evaluate the effects of a treatment more accurately.

These definitions also draw on distinguishing marks that belong to quite different categories, namely natural kinds, roles, and being the subject of a procedure. Such categorial distinctions should be honored in a well-constructed ontology.

2.4. *Meta-Types and ‘Other’*

The NCIT is also deficient with regard to explicitness and precision. Like the CAT, the NCIT contains the entry ‘Other’. This is a subtype of ‘General_Modifier’ (which is a subtype of ‘Qualifier’ that, in turn, is a subtype of ‘Properties_and_Attributes’) and is defined as ‘Different than the one(s) previously specified or mentioned’. In all, there are approximately 80 *other*-involving entries in NCIT including for example: ‘Carcinoma, Other, of the Mouse Pulmonary System’.

Another trait the NCIT shares with CAT is that of including meta-types (types that are dependent on the classification of which they are a part) alongside types within its hierarchy. For example, NCIT contains the type ‘NCI-Thesaurus_Property’, which is a subtype of ‘Property’ and is defined as a ‘specific terminology property present in the NCI Thesaurus’. Meta-types even occur at the top-node level of the NCIT: its top-level features the heading ‘Retired_Concept’, defined as: a ‘Concept [that] has been retired, and should not be used except to deal with old data’. This entry clearly mixes properties of the term with properties of the entities to which the term refers (compare Frege, 1884, § 53, and 1892, 192-205). Although it is undoubtedly useful to have a record of a term’s properties, these properties should not be dealt with as if they were characteristics that a thing must have in order to instantiate a certain universal.

3. *Restrictive Conditions for Classifications*

In criticizing Borges’s CAT and the NCIT, I have been guided by a vision of an ideal classification. According to this ideal, a classification consists of classes that are jointly exhaustive and pairwise disjoint (JEPD) and constructed out of ontologically well-founded distinguishing characteris-

tics. There are a number of reasons why real-life classifications deviate from this ideal image.

A first group of limitations on classification derives from the domain to be classified. Particularly in the case of biological kinds, we have the problem that there is a large number of, for example, animal or plant or protein kinds which have not yet been scientifically described or even discovered. In addition, new genetic methods are enabling scientists to discover distinctions between kinds that are not available to traditional phenotype-based methods. The sheer number of kinds guarantees that biologists will have their work cut out for them for the foreseeable future. The number of animal kinds, alone, is estimated at approximately 30 million. There may be areas, such as human anatomy, that are close to being perfectly understood. But other areas are subject to constant growth in knowledge, such as zoology, botany, and especially genetics, which, because of the amount of available data, would hardly be possible to organize without the support of computers. Above all, however, we must bear in mind the likelihood of new species being discovered; not least because new species are constantly coming into existence. Such considerations, relating specifically to the domain to be classified, pose strict limitations on the exhaustiveness of a classification system. Some domains pose more principled problems for classification. Since, for example, bacterial genes can be switched from one bacterium to another and, because of the high rate of bacterial reproduction, can undergo rapid change, it is particularly difficult to distinguish stable species and kinds of bacteria.

A second group of limitations on classification derives from the technical side of the creation and application of classification systems. It does not matter whether we are dealing with a traditional, printed format, or a computer database; in either case, storage space is finite. Should computer programs be used for automated reasoning with the data contained within a classification, we have the problem of computability in addition to the problem of storage space. The time required for computation grows with the total number of classes, and with the number of inter-class relations with which a program must deal. Also, depending upon the programming language and its underlying logic and expressive power, there is the danger that a given task might not even be computable at all.

In addition to limitations posed by the domain of classification and by hardware and software, there are limitations posed by the human user. For

while it is becoming ever easier and cheaper to extend the storage space on computers, the cognitive abilities of their human users have narrow limits. Human archivists and librarians are advised to use no more than one thousand systematically ordered key words (approximately) to index books or documents (Gaus, 2003, 93-94). Computers can, of course, use many more terms than this; the *NCI Thesaurus* with its 36,000 words is not a particularly large terminology database. As early as 2001, for example, the *Unified Medical Language System* (UMLS) encompassed 1.9 million expressions with more than 800,000 distinct meanings (see Dugas and Schmidt, 2003). But it is human curators who construct and maintain such artifacts, just as it is humans who later use them. The curators are experts who often specialize in the development of this particular kind of knowledge representation. But when, say, a general practitioner uses a certain classification as a diagnostic coding system in the process of billing, we have to ask how many diagnostic codes we can reasonably expect to be used in everyday practice.

Thus, there are several explanations for the deviation of real-life classifications from our envisioned classificatory ideal, and the main reason is that there are certain trade-offs between our various goals. If we want a complete representation of a given scientific domain, this might be far from a system that is easily comprehensible for a human user. If achieving completeness means to amass large amounts of data and to encode many relations between classes, we may also run into problems of computability. If, on the other end, we use simplifying types like *other* or *not otherwise specified*, we may run into trouble when updating the classification; for in the different versions *other* may have a quite different meaning and, thus, a different extension. But if we refrain from using other-types and simply give up the JEPD criterion, we lose a considerable amount of inferential strength. For, then, we no longer know that an entity that belongs to a supertype also belongs to one of the respective subtypes, and so on.

4. *Reference Ontologies: A Possible Solution*

A recent suggestion to solve this dilemma is based on a clear division of labor. We simply need two kinds of systems: reference ontologies and application ontologies. *Reference ontologies* should be developed without any regard to the problem of storage and the processing time, and they

should represent, at any given time, the state of knowledge of the respective scientific discipline from which they derive (see OBO, 2006):

A reference ontology is analogous to a scientific theory; it has a unified subject-matter, which consists of entities existing independently of the ontology, and it seeks to optimize descriptive or representational adequacy to this subject matter to the maximal degree that is compatible with the constraints of formal rigor and computational usefulness. Because a reference ontology is analogous to a scientific theory, it consists of representations of biological reality which are correct when viewed in light of our current understanding of reality (and thus it should be subjected to updating in light of scientific advance).

An *application ontology*, on the other hand, is analogous to a technical artifact like a computer program. Up to now, it was customary to build new ontologies from scratch for each new kind of application. This causes much trouble for anyone who wants to exchange or compare data among these different systems. It is better to use an already-existing reference ontology, from which we can derive the application ontology through a choice or combination of types from the reference ontology that are appropriate to the respective aim of the application ontology. Then, several such application ontologies can be mapped to each other through their respective reference to a common reference ontology.

While the task of maximally adequate representation of reality is transferred to the reference ontology, the application ontologies are constructed in light of the limitations posed by storage space, processing time, and the needs of the human users. While reference ontologies care about scientific virtues like completeness and precision, application ontologies care about engineering virtues such as efficiency and economic use of resources. The scientists of the OBO Foundry (see Smith, *et al.*, 2007) regard this as decisive progress:

The methodology of developing application ontologies always against the background of a formally robust reference ontology framework, and of ensuring updating of application ontologies in light of updating of the reference ontology basis, can both counteract these tendencies toward ontology proliferation and ensure the interoperability of application ontologies constructed in its terms. (OBO, 2006)

5. Exotic Thinking or Unfit Tool?

Some philosophers have joined with Foucault in claiming that Borges's CAT possesses a certain exotic charm (Foucault, 1970; see also Jullien, 1990). I have shown that CAT is charming indeed, in that it can illustrate a wide range of possible mistakes in constructing taxonomies. CAT is, of course, literature and not science. As a contribution to science, it would not be evidence of exotic thinking, but rather of impractical thinking. For its part, the NCIT is not a piece of literature but is intended to be a piece of science. And it is, we believe, an example of very impractical thinking. In fact, the National Cancer Institute which maintains the NCIT is indeed itself dissatisfied with the present state of its thesaurus and its purported exotic charm, and is taking steps to improve it. As I have shown, such emendation is an excellent proof that technical applications can be helped by being built on foundations laid by philosophy.

Chapter 8: Categories:

The Top-Level Ontology

Ludger Jansen

The task of ontology is to represent reality or, rather, to support the sciences in their representation of reality. In the last chapter, the reader became acquainted with an important means of doing so, namely: the technique of classification. But, in any classification, what are the very first kinds? What should the top level look like? In this chapter, I attempt to answer these questions. First, I review some suggestions for top-level ontologies with the help of the criteria established in Chapter 7 (section 1). From the point of view of the philosophical tradition of ontology, the question of a top-level ontology is tantamount to the question of the most basic categories. In order to develop some alternative suggestions, the nature of categories must first be addressed. To this end, I appeal to the philosopher whose ideas are pivotal in influencing our current understanding of ontology: Aristotle (section 2). Starting from Aristotle's list of categories (section 3), I go on to discuss three dichotomies which I recommend as candidates for the seminal principles of a top-level ontology, namely: dependent versus independent entities (section 4), continuants versus occurrents (section 5), and universals versus particulars (section 6). Finally, I discuss some categories of more complex entities like states of affairs, sets, and natural classes (section 7).

1. *SUMO, CYC & Co.*

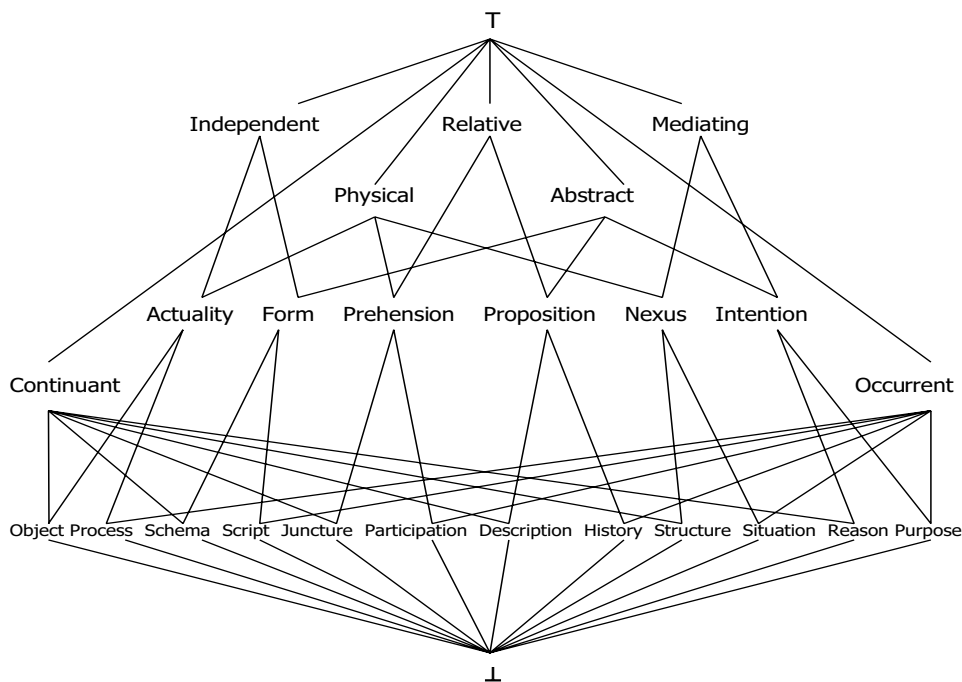
What should an ontology look like at the highest level? What are the most general classes of all classifications? Authors in the fields of informatics and knowledge representation have offered various suggestions. Some of the best known are:

- the OpenCyc Upper Ontology: the open-source version of the Cyc technology, developed by the Texas-based ontology firm Cycorp, which is supposedly the largest implementation of general knowledge inside a computer for purposes of common-sense reasoning;²³

²³ See *Cyc*, as of August 8, 2006: 'OpenCyc is the open source version of the Cyc technology, the world's largest and most complete general knowledge base and commonsense reasoning engine'.

- SUMO, the Suggested Upper Merged Ontology, which developed from an open-source project bringing together freely available, non-commercial ontologies into a common system; together with its various domain ontologies SUMO, supposedly, is currently the largest publicly accessible ontology;²⁴
- the Sowa Diamond (see Figure 1), representing in graphic form the top-level ontology suggested by John Sowa, which forms twelve categories by means of two dichotomies and a trichotomy in a lattice-like array (see Figure 1);²⁵
- BFO, Basic Formal Ontology, developed by the Institute for Formal Ontology and Medical Information Science (IFOMIS), and which exists in three versions (OWL DL, First-Order Logic, and OBO format).²⁶

Figure 1: *The Sowa Diamond*²⁷



²⁴ See *Ontologyportal*, August 8, 2006: ‘The Suggested Upper Merged Ontology (SUMO) and its domain ontologies form the largest formal public ontology in existence today’.

²⁵ Compare Sowa, 2000, 2001.

²⁶ See *BFO*; Grenon, *et al.*, 2004; Grenon and Smith, 2004; Grenon, 2003.

²⁷ Source: John F. Sowa. ‘Top-level Categories’, <http://users.bestweb.net/~sowa/ontology/toplevel.htm> (August 8, 2006).

In the following, I am going to compare OpenCyc to the quality criteria for classifications expounded in the last chapter. The suggestion for an Aristotelian-inspired top-level ontology, which will be developed in what follows, corresponds to the most basic traits of BFO, building on the three dichotomies between independent and dependent entities, continuants and occurrents, and universals and particulars. Over the course of developing these suggestions, it will become clear where the Sowa Diamond needs to be repolished (section 8).

In contrast to the completely symmetrical Sowa Diamond, the top level of the OpenCyc Upper Ontology is a complicated ('tangled') conglomerate. The graphic representation of this classification system in Figure 2 gives us an impression of this.

Against the background of the criteria for classifications addressed in Chapter 7, issues with the highest dichotomy in this diagram become immediately apparent. Why should we divide the class *thing* into the two subclasses of *Individual* and *PartiallyIntangible*? These two classes are neither jointly exhaustive nor pairwise disjoint. The latter, it seems, was introduced to have a place for persons, who putatively embody both tangible and intangible (mind-related) aspects. OpenCyc quite clearly admits of multiple inheritance, which manifests itself in diamond-like structures in the diagram. The reader will notice the combined subclass of *PartiallyIntangibleIndividual* at the level below these two classes. The two classes mentioned do not exhaust the class of *Thing*. Non-individuals (that is, the universals) do not appear as such in the diagram. The categories placed in opposition to the Intangible, namely, *PartiallyIntangible* and *TangibleIndividual*, do not appear in the diagram until four levels later.

Further, the diagram does not distinguish sufficiently between classificatory differences (such as *PartiallyTangible*) and the classes thereby engendered (such as *TangibleThing*). When we read the connective lines in the sense of the *is_a* relation, as we should be able to do in a classification system, then what results is grammatical nonsense: *TangibleThing is_a PartiallyTangible*. The subsumption relation *is_a* does not find application here. An ordinary predicative structure would be much more appropriate here, as in: *TangibleThing is PartiallyTangible*.

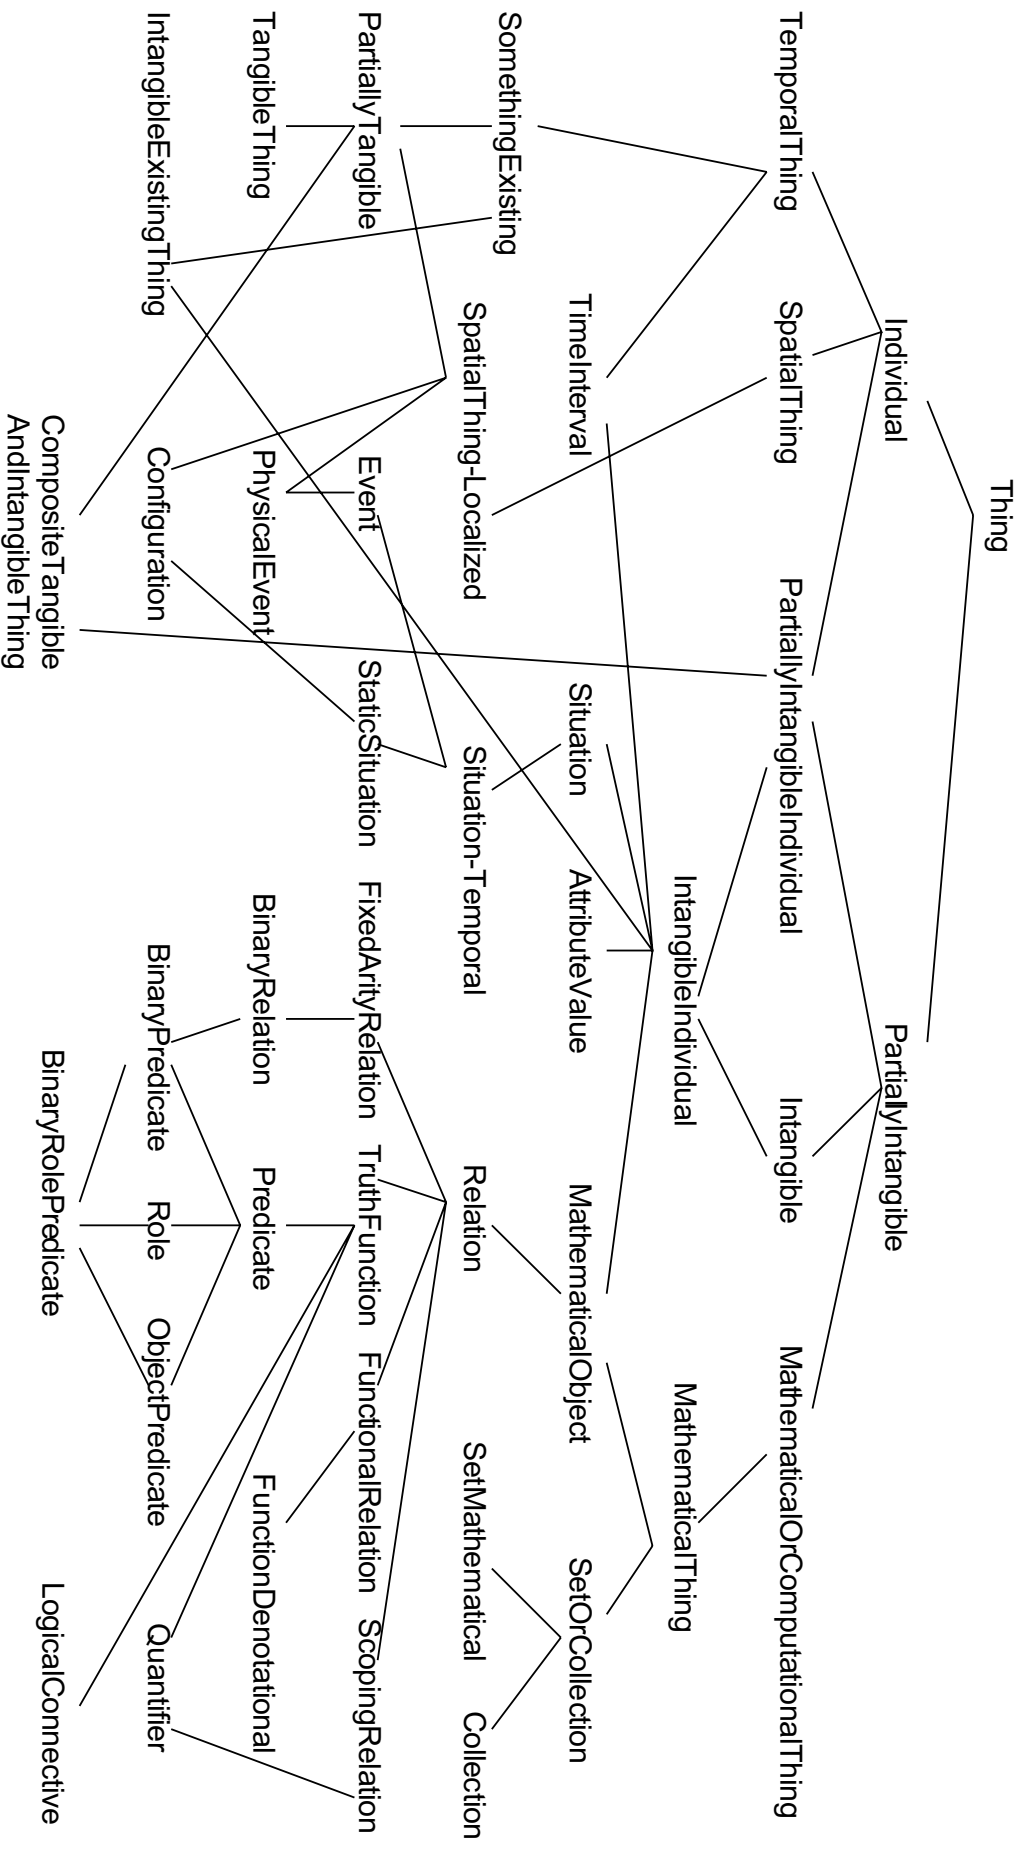


Figure 2: *The OpenCyc Upper Ontology*.
 Source: <http://www.cyc.com/cycdoc/upperont-diagram.html>

It is surprising that, apart from these problems, the property of tangibility is given such a prominent position in the first place. Attributes such as spatiotemporality or materiality seem to be much more basic and, also, better understood. Like many predicates expressing dispositions, ‘tangible’ is an extremely ambiguous term. God, an electron, the Milky Way, the Earth’s gravitational field, the country of Germany, Beethoven’s Fifth Symphony, a sound wave, meanings, neighborliness, freedom, a football game, an hour, yesterday’s snow, the exponential function, a computer program, my conception of the moon, and a stone enclosed in epoxide resin are all intangible, but for very different reasons. These reasons indicate aspects of these things that would make better traits on which to base an ontological classification.

OpenCyc’s subsumption relations are also problematic with respect to details. *TimeInterval* is surely a *TemporalThing*, but is it an *Individual*? In any case, not in the sense of indivisibility (or more precisely: the inability to be divided into two things of the same kind as the thing divided), for every time interval can be divided into parts which are themselves time intervals. On the other hand, *SituationTemporal* does indeed seem to be a *TemporalThing*. The class *Relations* is subsumed under *MathematicalObject*. Yet, my being in love with someone, being somebody’s neighbor, and being an employee are all relations, but they are not mathematical objects. Similarly, my stamp collection is a *Collection*, but it is by no means a *MathematicalObject*, and it is tangible all over; thus, in no way is it a *PartiallyIntangible* thing.

No ontologically apt classification principles can be found in the diagram’s ‘or’ expressions *MathematicalOrComputationalThing* and *SetOrCollection*, for there seems to be no good reason to treat the result of combining two universals by means of an ‘or’ relation as constituting a universal in its own right (Armstrong, 1978, II, 19-23). The class *SomethingExisting* is also strange – do the other classes comprehend entities that do not exist? Here the property of existence is wrongly being treated as a characteristic of things (see Frege, 1884, 53, and 1892, 192-205). The highly varied division of relations is ultimately based, mainly, on logical considerations; but these are entirely independent of the ontology of relations (see Jansen, 2006).

All of these are good reasons to work towards a more unified and consistent form for the uppermost levels of classification systems appropriate for ontologies. In what follows, such a unified form will be

developed drawing upon one of the oldest suggestions for such a top-level ontology, namely: Aristotle's *Categories*.

2. *What are Categories?*

As far as we know, Aristotle was the first to use the Greek word *kategoria* as a technical term in the context of philosophy. Originally, the noun *kategoria* and its corresponding verb, *katêgorein*, belonged to legal discourse. There, *kategoria* means the accusation in front of the judge, and *katêgorein* means to accuse someone. Probably because an accusation asserts something of someone, the verb can also mean 'make known' or 'assert', and was used in this way by Plato.²⁸ Aristotle uses the active verb phrase *katêgorein ti tinos* in the sense of 'to assert something about something', but even more often he uses the passive *katêgoreisthai ti tinos* or *katêgoreisthai ti kata tinos* in the sense of 'is said of something'. Aristotle uses the noun *kategoria* as the technical term for predication or for the predicate itself. In addition, he uses the plural of the noun in the sortal sense of 'kinds of predicates/of predication', and it is only in this usage that the Greek word *kategoria* can be translated into English as *category* (Jansen, 2006).

We have evidence that Aristotle's conception of the categories developed in three phases. First, as in *Topics* I 9, the distinction of different categories was only meant as a classification of predicates. In this first phase, the categories served as aids for finding arguments and for avoiding or discovering false inferences; thus, they had their place in the theory of argumentation. The second phase is represented in Aristotle's *Categories*. There the division of categories encompasses, not only predicate terms, but also subject terms. In this phase, terms denoting so-called primary substances, i.e. proper names such as 'Socrates' or 'Brunhilde', fall under the first category of substance, although they can function only as the subject of predication but never as predicates (*Categories* 5, 3a 36-37). This represents a step away from the theory of argumentation in the direction of ontology. In the third phase, which finds its expression in the *Metaphysics*, we find Aristotle's famous observation that 'to be' and 'a being' are used in as many different ways as there are categories (*Metaphysics* V 7, 1017a 22-23). Here, the division into separate categories

²⁸ See e.g. Plato, *Theaetetus* 208b; *Phaedrus* 73b. *Theaetetus* 167a links both meanings with each other.

became a full-fledged part of one of the most important of Aristotle's ontological teachings.

Aristotle's theory of categories was the subject of much dispute in antiquity, and has been interpreted in a variety of ways in the history of philosophy. Partly, this has to do with the fact that category theory had many different facets, even in the works of Aristotle himself. This came about because either Aristotle subjected his ideas to further development, or highlighted different aspects when presenting his theory. We can distinguish four prototypical interpretations (which often appear in combination), according to whether the categories classify (1) subject and predicate *terms* and the associated *meanings*, (2) *beings*, (3) mental or extra-mental *concepts*, or (4) *meanings of the copula 'is'*.²⁹ Here, we can draw on what was certainly the main conception of the late Aristotle, namely: that of the categories as the highest species of beings.

3. Aristotle's Ten Categories

In *Topics* I 9, Aristotle says explicitly that there are ten categories, which he then proceeds to delineate. A list of ten categories can also be found in the *Categories* (see Figure 3). Aristotle names many of his categories with the interrogative expressions that one would use to ask questions whose answers would make reference to entities in the respective categories. Many of the current names for these categories have their origins in the corresponding Latin interrogative expressions.

Figure 3: *Different Terms for Aristotle's Categories*

ARISTOTLE'S TERM	ENGLISH TRANSLATION	LATIN TERM	MODERN TERMS
<i>ti esti, ousia</i>	What is it?, essence	quod est, quiditas, essentia	essence
<i>poson</i>	How much?	quantum, quantitas	quantum, quantity
<i>poion</i>	How is it?	quale, qualitas	quality
<i>pros ti</i>	Related to what?	relativum	relative, relation
<i>pou</i>	Where?	ubi	place
<i>pote</i>	When?	quando	time
<i>keisthein</i>	lying, being situated	situ	position, posture
<i>echein</i>	having	habitus	
<i>poiein</i>	doing	agere	
<i>paschein</i>	suffering	pati	

²⁹ See Bonitz, 1853; Ebert, 1985; Kahn, 1978; Oehler, 1986.

Kant accused Aristotle of choosing his categories in a rhapsodic manner. In this unsystematic way, Aristotle could never be certain that his list of categories was complete (Kant, 1781, A 81 = B 106-107). Later Aristotelians, such as Thomas Aquinas³⁰ or Franz Brentano (1862; see also Simons, 1992), undertook the task of constructing a system that yields the Aristotelian categories, in the precise order in which they are named and discussed in the *Categories*.³¹ We can assume that Aristotle himself constructed his list of categories indeed in an unprincipled way, as Kant suspected, for he seems to have proceeded simply on the basis of his experience in dialectical exercises and philosophical discussions.

This might explain the disparity of Aristotle's list of categories, since the elements in his list are not at all of the same standing. There are two important ways in which Aristotle's categories fall into disparate groups, which I will discuss in due course: They encompass dependent as well as independent entities (section 3), and continuants as well as occurrents (section 4). These are already two of the ontological dichotomies that can be used as the seminal principles of the top-level ontology. Following these, I will introduce a third dichotomy that is orthogonal to the other two: the distinction between universals and particulars (section 5).

4. *Dependent and Independent Entities*

In the *Categories*, Aristotle distinguishes between primary substance (*protê ousia*), that is, a substantial particular, and secondary substance (*deutera ousia*), a species of substantial particulars. Of these two, Aristotle accords special ontological status to the individual substances. Everything else is either predicated of these individual substances, or is *in them* as something underlying them (*Categories* 5. 2a 34-35; 2b 3-5; 2b 15-17). In later texts as well, Aristotle accords this first category of individual substance a special importance with respect to the other categories, which are also called 'affections of the *ousia*'.³² Aristotle is quite clear that his ten categories are not to be viewed as equals; rather, the individual substances

³⁰ See Aquinas, *In Physicorum Aristotelis expositio* III, lectio 5, Nr. 322 [15] and *In Metaphysicorum Aristotelis expositio* V, lectio 9, Nr. 891-892.

³¹ See Jansen 2007 for a new suggestion of a hierarchy of Aristotle's categories along the lines suggested here.

³² *Metaphysics* IV 2, 1003b6: *ousiai – pathê ousias*; see also *Metaphysics* XIV 2, 1089b 23: *ousiai – pathê – pros ti*.

are presupposed by the other categories. From Aristotle's perspective, it is this fact that made the unity of ontology possible (*Metaphysics* IV 2).

Customarily, the dependent categories are called *accidents* and are placed in opposition to substances. A traditional criterion for the opposition of substances and accidents can be found in the second chapter of the *Categories*: qualities and quantities are in a substance, while substances are not in a substance but, rather, are identical with one. But it is not entirely clear how this 'being in something else' is to be understood; for a heart is in a body and a tapeworm is in a host. This could not be the type of 'being in something else' that Aristotle meant. Aristotle explicitly excludes 'being-in' in the sense in which a part is in a whole as the heart is in the body. But a parasite such as a tapeworm is not a part of its host.

The criterion of ontological dependence helps to solve this problem. The tapeworm could leave its host and move into another host. A grin, a certain height, or a certain color could not leave their bearers in this way and continue to exist. It is not possible for the Cheshire Cat to disappear and leave its grin behind.³³ The height of a tree cannot continue to exist when the tree is destroyed. The color of a test tube cannot remain in a room when the test tube is taken out of the room. The grin, the height, and the color are dependent for their existence upon a bearer, a substance which has this grin, this height, or this color, among its properties. They cannot migrate from this substance to another: if Alice were to grin instead of the Cheshire Cat, then it would be a *new* grin.

Let us summarize this thought. Substances do not need the entities of other categories in order to exist, whereas the entities of other categories require entities from the first category for their existence. For this reason, substances are called *ontologically independent* entities, where accidents are said to be *ontologically dependent*. More precisely: substances are ontologically independent of accidents, while accidents are ontologically dependent upon substances. The notion of ontological dependence can be formally captured through a counterfactual criterion:

Def. (6.1) An entity x is ontologically dependent upon an entity y if x could not exist if y did not exist.

³³ Lewis Carroll, *Alice's Adventures in Wonderland*, Chapter 6: 'I've often seen a cat without a grin, thought Alice; but a grin without a cat! It's the most curious thing I ever saw in my life!' (Carroll, 1965, 67).

For substances and their accidents it holds that: if s is a substance and a is one of s 's accidents, then a cannot exist unless s exists. Because a inheres in s , a is ontologically dependent upon s . On the other hand, however, not all of those things that are ontologically dependent on other entities inhere in those entities. A relational event such as a kiss or a hit are ontologically dependent upon their relata, but they do not of inhere in any of their relata; rather, they inhere in the totality which these relata form.

It is possible for two entities to be mutually ontologically dependent. Someone can only be a patient when there is a doctor treating him, and there can only be an active doctor when there is also a patient. Now, being a doctor is not dependent upon the existence of a particular individual patient; any patient, at all, would be sufficient. By the same token, the existence of patients does not end when a single individual doctor ceases to exist. Only if there are no more doctors whatsoever can there be no more patients. Doctors and patients are thus *generically* dependent upon one another. We can define generic dependence as:

(Def. 6.2) Being F is generically dependent upon being G if nothing can be F unless something is G .

On this definition, generic ontological dependence is a relation between universals.

We had defined ontological dependence in such a way that it is a relation that could obtain, in principle, between entities in any category; thus ontological dependence can also obtain between universals, according to the following definition:

(Def. 6.3) A universal F is ontologically dependent upon a universal G if the universal F cannot exist unless universal G exists.

The best criterion for determining whether the existence of a universal F presupposes the universal G , is to ask whether F could exist if nothing at all is G , and this is precisely the definition of generic dependence. Hence, there is no difference between the generic dependence of being F on being G , and the ontological dependence of the universal F on the universal G .

The group of accidents can be further divided into relational and non-relational entities. Relational entities are those that are ontologically dependent on multiple bearers, while non-relational entities are those that are ontologically dependent upon one bearer only (see Jansen, 2006; Smith and Ceusters, 2007).

5. *Continuants and Occurrents*

There is another way in which Aristotle's list of categories is not uniform. Two of the Aristotelian categories, those of *action* and *passion*, differ in an important way from the others. Whereas a substance such as a bacterium, a quantity such as a length of 20 meters, or a quality such as red, exist *in toto* at every point in time at which they exist at all, the existence of actions and passions is spread out over the course of some time interval. Whenever we encounter a bacterium, we encounter the *whole* bacterium at each point in time over the course of the bacterium's life. The process by which a bacterium reproduces, by contrast, or a process such as healing, take place within time and are manifested over a time span. The process of reproduction has a beginning and an end; it is composed of various phases that follow one another in time. These entities, reproduction and healing, have temporal parts. By contrast, the bacterium has spatial parts – for example, a nucleus, a membrane, and a cytoplasm – which exist at one and the same time.

Hence, we see that there are two kinds of entities that stand in intimate relation to one another, namely: (1) an organism and (2) its *life* or *history* (which might be documented in a patient record). The organism itself is present as a whole at every point of its existence, while the life of the organism is spread out over multiple points in time. In the former case we are dealing with entities which continue to exist through time, which we call *continuants*. In the latter case, by contrast, there is no point of its existence at which the entity is wholly present. It unfolds in time, that is, it has temporal stages or phases. The latter are not identical with one another, but are rather various different parts of the temporal entity. These are things that occur in time, and for this reason are called *occurrents*.

The words 'continuant' and 'occurrent' can be traced back to the Cambridge logician William Johnson (the teacher of Bertrand Russell). Johnson defines 'continuant' as 'that which continues to exist while its states or relations may be changing' (1921, 199). More recently, David Lewis (1986, 202) drew a similar distinction between *endurers* and *perdurers*:

Something *perdures* iff it persists by having different temporal parts, or stages, at different times, though no one part of it is wholly present at more than one time; whereas it *endures* iff it persists by being wholly present at more than one time.

Distinguishing between these two modes of existence is often seen as marking a distinction between two alternative, and competing, theories of the diachronic behavior of the same entities. David Lewis, for example, claimed that *all* entities must be seen as four-dimensional perdurers (thus as occurrents).³⁴ Here, instead, we will argue that Socrates and his walking exhibit two very different modes of existence. While the walking is clearly an occurrent, Socrates himself is no less clearly a three-dimensional continuant. Hence, there are two kinds of entities which demand distinct theories to account for their diachronic behavior. We need both continuants and occurrents in order to represent reality accurately.

But the opposition between continuants and occurrents does not present an exhaustive classification of all entities. For this opposition appears only with those entities whose existence, in fact, is extended over multiple points in time. There are at least two problem cases which this distinction does not encompass, namely, instantaneously existing qualities and quantities (see Johansson, 2005), and points in time themselves. It is trivially true that a point in time exists only at one point in time, that is, at itself. And, in processes of growth and change, it is possible for instantaneously existing quantitative and qualitative individuals to be substituted for each other. If a ball grows continuously at a constant rate during this growth process there are no two points at which the ball has the same weight. If a surface changes its color continuously from, say, blue to red, at no two points in time is this surface the same color. Since the existence of these instantaneous qualities and quantities does not extend over multiple points in time, it would seem to follow that there are qualities and quantities which fall under the category of continuant, as well as those which do not. In the same way, time intervals would belong to the category of occurrent, but points in time would not. This does not make for a particularly elegant theory. So, we will modify these categories slightly, in order to integrate these homeless entities.

If we picture the world at any single point in time, we will discover people, animals, artifacts, colors, sizes, and relations in our picture. But changes, processes, and events that are taking place at that point in time will not be visible in the picture. In order to represent these, we need a sequence of pictures instead of a single picture; we need a film. In order to obtain a complete picture of our ever changing world, we thus need two kinds of representation.

³⁴ For an overview of this discussion, see e.g. Lowe, 2002, 49-58.

On the one hand, we need snapshots of the world at particular points in time, which capture the continuants. Let us call such snapshots *SNAP ontologies* (following Grenon and Smith, 2004). Included among SNAP entities are substances, quantities, qualities, relations, as well as the boundaries of substances, collections of substances, places such as niches and holes, and spatial regions such as points, lines, surfaces, and volumes. Over and above to the traditional category of continuants, SNAP ontologies comprise also the merely instantaneously existing instances of qualities and quantities which would otherwise be ontologically homeless.

On the other hand, we need a representation of change, something like a film which represents entire time spans. We will call these *SPAN ontologies* (after Grenon and Smith, 2004). Included among SPAN entities are happenings such as processes and events, temporal regions such as time intervals with time points as their boundaries, as well as spatiotemporal regions. In Chapter 12 we will discuss happenings, the specific elements of SPAN ontologies. Time points, in spite of their lack of temporal extension, belong to the SPAN ontology and not to the SNAP ontology. A single SNAP ontology, which represents the world at a given point in time, is linked to this time point as to its date, but does not contain this time point as one of the entities in its coverage domain.

6. *Universals and Particulars*

In addition to the two ontological dichotomies already discussed – independent vs. dependent entities, continuants vs. occurrents – there is also a third: that between universals and particulars. Since this distinction cuts straight through all of the Aristotelian categories, we can call it transcategorical.³⁵ This third distinction is also given systematic treatment in Aristotle's *Categories*. In the second chapter, he distinguishes between what can and what cannot be predicated of another entity. Predication requires an aspect of generality. Particulars, such as Socrates or my height, cannot be attributed to other entities. Sentences that contain as predicates the expressions 'is Cicero' or 'is my height' are not predications in the technical sense, but rather identity claims like 'Tully is Cicero' or 'Five feet is my height'. A general expression such as 'human' can appear both

³⁵ See Lowe, 2006, 21: 'The terms 'particular' and 'universal' themselves, we may say, do not strictly denote categories, however, because they are transcategorical, applying as they do to entities belonging to different basic categories'.

as the subject and as the predicate of predicative assertions, as in ‘A human is a vertebrate’, and ‘Cicero is a human’.

Taken together with the distinction between inhering and non-inhering entities, this yields a fourfold distinction of entities, the so-called *ontological square* (represented in Figure 4).³⁶ Many ontologists accept only a selection of the fields of this ontological square. David Armstrong, for example, tries to manage with fields I and IV only, namely, particular substances and property universals (Armstrong, 1978 and 1997). Ontologists who see First-Order Logic on its standard reading as a tool for ontology arrive at the same result. The particulars correspond on this account to the individual constants (‘*a*’, ‘*b*’, ‘*c*’ ...), and the property universals correspond to the predicate variables (‘*F*’, ‘*G*’, ‘*R*’ ...). The view that the formula ‘*F(a)*’ is the key to ontology – that such formulae, along with relational expressions such as ‘*R(a, b)*’, in effect, form a mirror of reality – has been dubbed *fantology* by Smith (2005a).

Those philosophers who are prepared to allow events into their ontologies, such as Donald Davidson (1980), also accept continuants, which intimately resemble entities in field II. Russell, by contrast, wanted to completely eliminate the level of individuals, and to satisfy himself with fields III and IV,³⁷ most likely having been influenced by Leibniz’s theory of individual concepts.³⁸ Nominalist philosophers, by contrast, accept only entities from the two lower fields, I and II. Some philosophers even try to make do with only one of these two categories. For example, the individual accidents in field II are the only basic entities for *tropists*; they call these abstract particulars or tropes,³⁹ and see individual substances such as you and me as more or less loosely connected bundles of such tropes.

³⁶ See Smith, 2003a. On the history of such diagrams see Angelelli, 1967, 12; see also Wachter, 2000, 149. One of the most important contemporary representatives of a four-category ontology is E. J. Lowe; see in particular 2006.

³⁷ See e.g. Russell, 1940, ch. 6; and 1948, Part II, ch. 3 und Part IV ch. 8; 1959, ch.9. For a similar position see Hochberg 1965, 1966, and 1969.

³⁸ Russell (1948) attributes this conception explicitly to Leibniz. See also Armstrong, 1978, I 89: ‘[...] while the influence of Leibniz to Russell is clear, it is less clear that Leibniz held this theory of the nature of particulars

³⁹ Two classic presentations of this position can be found in Williams, 1953, and Campbell, 1990. See also Macdonald, 1998 and Trettin, 2000.

Figure 4: Aristotle's Ontological Square

	<i>Substantial</i> Not in a subject	<i>accidental, non-substantial</i> In a subject
<i>universal, general</i> Predicated of a subject	III. <i>substance universals</i> Human being Horse	IV. <i>non-substance universals</i> Being white Knowing
<i>Individual</i> Not Predicated of a subject	I. <i>individual substances</i> This human being This horse	II. <i>individual accidents</i> This individual whiteness This individual knowing

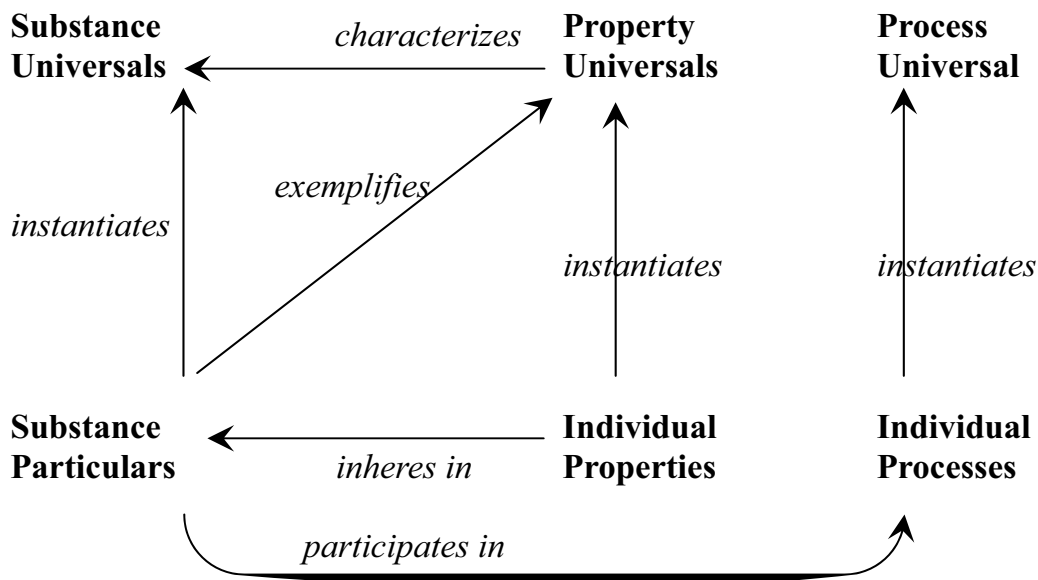
Aristotle accepted all four cells of the ontological square, which he sees as, together, forming a transparent partition of reality. Thus, he reflects the commonsensical understanding of most people, according to which elements of all four fields exist. In daily life, we assume that George W. Bush (field I) exists as well as the species *elephant* (field III), the virtue of courage (field IV), and the individual white color of my skin, which ceases to exist at some time in summer, when my skin takes on a brown color instead (field II). Ontologists who want to get rid of one or more of these fields represent some kind of reductionist position. They must produce an alternative explanation for why we suppose in our everyday understanding that these things exist. They do this mainly through explaining our reference to entities in these fields as merely a roundabout way of talking about entities in other, more highly favored, fields.

There are some basic relations that obtain among entities in the four fields of the ontological square:

- Individual accidents *inhere in* individual substances.
- Non-substance universals *characterize* substance universals.
- Individual substances *instantiate* substance universals.
- Individual accidents *instantiate* accident universals.
- Individual substances *exemplify* accident universals.

A picture of the world which did not provide a special place for occurrents would be incomplete. There are of course important relations that obtain between occurrents and continuants, for there are individual substances which take part in individual processes and events. We can thus expand the ontological square to an *ontological sextet*, which can be illustrated in Figure 5 (Smith, 2005a). The relations of inherence, exemplification, instantiation, and participation govern the relations among the entities in these four fields. They are important formal-ontological relations; regardless of which area of reality we want to represent, we must take all of these relations into account.

Figure 5: *The Ontological Sextet and the Formal-ontological Relations*



7. Complex Entities

In addition to the categories we have discussed thus far, discussions take place among modern ontologists about complex entities such as states of affairs, sets, mereological sums, and classes.

States of affairs are all of those complex entities which can be described with a ‘that’ sentence. That the ball is round and that the cat is on the mat are two examples of states of affairs. Both are complexes of entities falling among the various categories which we have just discussed. That a person is sick is a complex composed of a substance (this person), and a certain quality or disposition (sickness). The state of affairs that a certain molecule is attached to a receptor is composed of: a substance (the molecule), a part

of a substance (the receptor), and the two-place relation of being attached. States of affairs can have other states of affairs as components. The state of affairs that the doctor has discovered that her patient has the flu is composed of the doctor, the intentional two-place relation of having discovered, and of the state of affairs that the patient has the flu. The thesis that all states of affairs are complex, or composite, entities seems to be called into question by expressions such as ‘that it rains’, which are constructed from impersonal pronouns such as ‘it’. For these expressions cannot be divided *linguistically* into a predicate, on the one hand, and a referring subject expression, on the other. But this does not mean that the entities for which they stand cannot be analyzed *ontologically*. The state of affairs that it is raining is clearly composed of raindrops moving from place to place; thus, it is composed of a collective of movements undergone by a multiplicity of raindrops.

Sets are well known from mathematics. Sets are collections of elements. We say that sets *contain* elements as their members. And we say that certain entities are (or are not) elements of certain sets. The relation *is an element of* is represented by the sign ‘ \in ’, while the relation *is not an element of* by the sign ‘ \notin ’. In addition, set theorists discuss a range of relations between sets such as the intersection, the union, the subset relations, and the relation of set-theoretical difference.⁴⁰ The intersection of two sets, for example, is the set – which may perhaps be empty – that contains as members exactly those entities which are members of both initial sets.

We can represent sets either extensionally, by listing their elements, or intensionally, by pointing to a feature common to all elements that is sufficient for set membership. Extensionally, sets usually are represented by means of lists whose elements are separated by commas and placed in closed parentheses. For example, the set of prime numbers less than 10 is {2, 3, 5, 7}. But {Aristotle, 2, my stethoscope} is a set as well; thus, sets can be built out of arbitrarily designated elements. To be sure we can represent sets intensionally, without such a list, simply by specifying the characteristics that the elements belonging to them share and that are sufficient for set membership. Examples of this sort of description of a set would be ‘the set of all patients at noon on the November 1, 2008 in Berlin’, or ‘the set of all such patients with a fever’. These sorts of descriptions are sometimes represented in the form: {x | x is a patient and

⁴⁰ For an overview see e.g. Bucher, 1998, Ch. 1.

has a fever}, which is read as ‘the set of all things x , such that: x is a patient and has a fever’. Additional examples of set descriptions are ‘ $\{x \mid x$ is round}’, and ‘ $\{x \mid x$ is red}’.

Sets are identical when they contain the same elements. The set description ‘ $\{2, 3, 5, 7\}$ ’ denotes the same set as the description ‘the set of prime numbers less than 10’, because each element contained in $\{2, 3, 5, 7\}$ is also contained in the set of prime numbers less than 10, and vice versa. The two sets, thus, are identical. From this criterion of identity, it follows that sets cannot survive the loss of any of their elements; the same set cannot have different elements at different points in time: different elements, different sets. From this criterion for set identity, it also follows that sets, in a certain sense, are timeless; hence, sets can include elements which exist at different times and at no times. They are also outside space (if the elements of a set move about in space the set is not affected in any way). It follows further that the order of the elements in a set is irrelevant. Thus:

$$\{a, b\} = \{b, a\}.$$

It also follows that repetitions of elements are irrelevant for set identity. Thus it holds that:

$$\{a, a\} = \{a\}.$$

In order to know whether $\{x \mid x$ is red} and $\{x \mid x$ is round} are the same sets, we must know what sorts of things are available in the world, or in some specially selected *universe of discourse*. If the world consisted merely in a red circle, a yellow triangle, and a blue square, then these two set descriptions would indeed denote the same set; that is, the set {red circle}. In the actual world, there are circles that are not red and, therefore, according to the criterion for set identity *in the actual world* these two sets are not identical. The criterion for set identity also entails that there are no two distinct empty sets.

Because sets are independent of space and time, they count as *abstract entities*. The curly brackets are a sort of mechanism of abstraction: we take the names of concrete entities, place brackets around them, and create a name for something abstract. From ‘Socrates’, the name of the flesh-and-blood Socrates who exists in space and time, we get ‘{Socrates}’; the name of an abstract entity, existing apart from space and time, that is the set

composed of Socrates as its only element. Sets containing only one element are called *singleton* sets. The empty set itself, which plays a prominent role especially in mathematical explorations of the implications of the axioms of set theory, is referred to by means of the symbol ' \emptyset '.

Sets can themselves be elements of other sets; and some sets have only sets as their members. There are also singletons of sets, and also the singleton of the empty set. Now this singleton can itself be an element of a set, for example of its singleton, and so forth. Thus the theory of sets sketched so far allows forming the singleton of the singleton of the singleton and so on of – the empty set. Hence it is possible to create potentially infinite structures out of nothing – more specifically, out of the empty set – and have these structures be isomorphic to the set of the natural numbers. Each of the following three rows fulfills the five Peano axioms for the natural numbers – only the interpretation of the neutral element 0 and the successor function are different:

0, 1, 2, 3, ...
 $\emptyset, \{\emptyset\}, \{\{\emptyset\}\}, \{\{\{\emptyset\}\}\}, \dots$
 $\emptyset, \{\emptyset\}, \{\emptyset, \{\emptyset\}\}, \{\emptyset, \{\emptyset, \{\emptyset, \{\emptyset\}\}\}\}, \dots$

Since the singleton of a concrete thing is an abstract entity, the singleton and its only element must be distinct from one another. This is 'the mystery of the singletons: what distinguishes a from $\{a\}$ '? (Simons, 2005, 145) The tricks that can be played with empty sets have induced some logicians and philosophers to seek an alternative to the set-theoretic view known as mereology (Simons, 1987; Ridder, 2002). *Mereological sums* are complexes which can be composed of various parts. My stomach, my sandwich, and the warmest corner of my office can comprise such a mereological sum. Just as with sets, there is virtually no limitation to the building of mereological sums. And just as with sets, many mereological sums (as in the example above) have a very artificial character. At any rate, very few mereological sums are natural wholes (though natural wholes such as organisms are among the most interesting of mereological sums). While sets are abstract entities even when composed of concrete elements, mereological sums composed of concrete elements are concrete things as well. Mereological sums exist in space and time, but only as long as all of their parts exist. A mereological sum does not survive the loss or destruction of one of its parts. Losing a part will result in *another* mereological sum. We speak of *proper* parts if we want to indicate that the

putative part is not identical with the whole. A non-proper part can, analogously to non-proper subsets, also be identical with the whole.

In many ontologies, part-whole relations are used as formal-ontological relations. The theory of granular partitions (Chapter 6) introduces an approach which attempts to blaze a third trail between set theory and mereology, in order to link the concreteness of mereological sums with the hierarchical nature of the element-of relation.

Where sets can have members of arbitrarily different sorts, we shall use ‘class’ in what follows to refer to collections of members which are in some sense constrained, as for example in: the class of mammals, the class of red things, the class of positively charged electrons. The category of class thus represents an attempt to do away with the arbitrary nature of set construction.⁴¹ Although ‘set’ and ‘class’ are often used as synonyms, we will use them to signify different things, as for example in SUMO, where

‘Set’ is the ordinary set-theoretic notion, and it subsumes ‘Class’, which, in turn, subsumes ‘Relation A’. ‘Class’ is understood as a ‘Set’ with a property or conjunction of properties that constitute the conditions for membership in the ‘Class’ (Niles and Pease 2001).

This also follows Smith, Kusnierczyk, Schober, and Ceusters (2006, 60) for whom ‘class’ signifies ‘a collection of all and only the particulars to which a given general term applies’.

When the general term connected to a class represents a universal, we can speak of a *natural class*: a natural class is the totality of instances of a universal. Whereas sets may be constructed by means of enumeration, natural classes require that there be universals of which they are the extension. Two natural classes are identical if they represent the same universal. Because not all general expressions correspond to universals, not all classes are natural classes. These non-natural classes are called ‘defined classes’, like for example: the class of diabetics in London on a certain day, or the class of hospitals in San Diego.

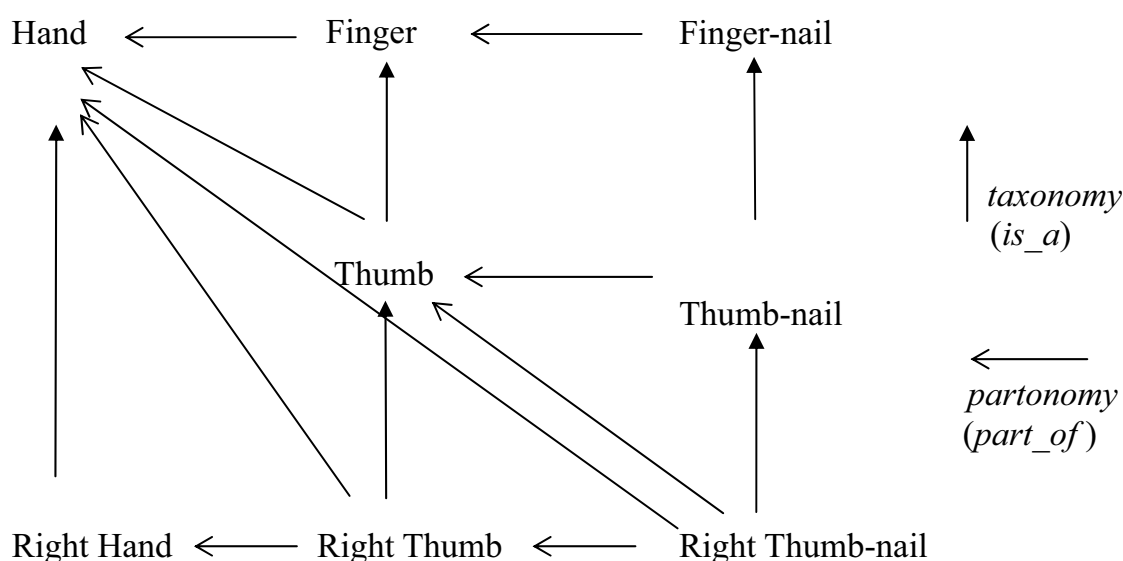
Not every set, on this view, corresponds to a class. For example, {Aristotle, 2, my stethoscope} is a set constructed through the listing of its elements. However, it does not correspond to a natural class, for it is not the extension of any universal; nor does it correspond to any class at all,

⁴¹ There are earlier attempts to link intensional elements with set theory; for example in Feibleman, 1974. The remarks presented here draw on Chapter 11 of this volume. See also Smith, *et al.*, 2005, Smith, 2005.

for there is no general expression (other than ‘element of that set’) under which precisely these three things fall. From a linguistic point of view we thus need, for the definition of a class, at least one general expression, whereas sets, such as the above example, can be denoted alone with proper names and definite descriptions.

Unlike set theory, class theory does not require us to know what things there are in the world in order to say that the class of red things and the class of round things are different from one another. And while there is only one empty set, there can be many different empty classes: for example, the class of all phlogiston, the class of all perpetual-motion machines, or the class of round squares. Since, however, they represent different universals, they are certainly different from one another. In addition, classes, but not sets, can survive the destruction or coming into existence of new instances; for sets are individuated by their elements, whereas natural classes are individuated by a universal which stays the same even as it has different instances at different times.

Figure 6: *A Combination of Taxonomy and Partonomy*⁴²



The result of dividing entities into classes is called a *classification*. Instead of speaking of a class we sometimes speak of a *taxon* (or, in the plural, of *taxa*, derived from the Greek word *tattein*, to place in order); we can speak, correspondingly, of a *taxonomy*. A taxonomy must be distinguished from a *partonomy*. While a classification or a taxonomy divides

⁴² From Zaiss *et al.*, 2005, 64.

a universal into species or kinds, a paronomy divides a whole into its parts. It is particularly interesting to combine a paronomy with a classification, which has been done in Figure 6.

8. *The Unpolished Edges of the Sowa Diamond*

We are now equipped to look more closely at the Sowa diamond. Sowa sees his ontology as a melting pot of the process ontology of Whitehead and the triadic category theory of Charles Sanders Peirce. In light of what we have already seen in this chapter, however, we can point to some things that have gone badly wrong in this melting pot. The systematic presentation of Sowa's ontology comprises a combination of three distinctions:

- a dichotomy between *Continuant* and *Occurrent*
- a dichotomy between *Physical* and *Abstract*
- a trichotomy (which Sowa attributes to Peirce) between *Independent*, *Relative*, and *Mediating*.

A first point of criticism could be the question whether the dichotomy *Physical* vs. *Abstract*, and the Peirce-inspired trichotomy, are in fact appropriate means of classification. I will not discuss this question here. These two dichotomies and the trichotomy, taken together, yield twelve combinatorial possibilities, which I would like to examine more closely.

Figure 7: *The Ten Central Categories of the Sowa Diamond*
from <http://users.bestweb.net/~sowa/ontology/toplevel.htm> (as of August 8, 2006)

	<i>Physical</i>		<i>Abstract</i>	
	<i>Continuant</i>	<i>Occurrent</i>	<i>Continuant</i>	<i>Occurrent</i>
<i>Independent</i>	Object	Process	Schema	Script
<i>Relative</i>	Juncture	Participation	Description	History
<i>Mediating</i>	Structure	Situation	Reason	Purpose

In contrast to Sowa, I do not find all of these combinations of di- and trichotomies well advised. For example, there are no abstract occurrents (see Guarino, 2001): what *occurs* is never abstract. Although there are universals that are instantiated by occurrents and only by occurrents, these universals are themselves not temporally extended entities and thus they

are not themselves occurrents (compare Chapter 12). To name an additional example: from our Aristotelian point of view, the category *Object* is the only one found among independent entities: all occurrents and all abstract entities are necessarily ontologically dependent entities.

Other combinatorial possibilities, like *Mediation* and *Participation*, seem to correspond more closely to what we would see as relations between categories than as categories in themselves. *Description* and *History*, by contrast, can both be understood as linguistic entities that are not distinguished ontologically, but rather by means of their objects. A description does not become an occurrent simply by being a description of an occurrent. Analogously, a *Purpose* does not become an occurrent simply because it aims at the realization of an occurrent (and even this does not hold for all purposes). Just as little is the general schema or recipe that describes how, e.g., an operation proceeds (what Sowa calls the *Script* of this event) thereby itself an occurrent. This is particularly clear when Sowa introduces a sheet of music and series of pictures on a roll of film as examples of scripts, as these exist in space and time and are thus, according to Sowa's own definition, physical entities and not abstract.

Sowa has designed his diamond in such a way that he characterizes the various options of his di- and trichotomies by means of axioms such that the central categories coming about through a combination of these options inherit the axioms of the options constituting them. Because of the problems just discussed it does not come as a surprise that this does not work. For example, Sowa characterizes occurrents *inter alia* as having sequential temporal phases and participants as spatial parts. The category *Reason*, which is characterized by Sowa as a *mediating abstract occurrent*, is meant to inherit these axioms. But reasons neither have temporal phases nor participants as spatial parts. Thus the principle of construction underlying the diamond cannot be held up.

An additional problem with Sowa's suggestion is that – notwithstanding its systematic outlook – it fails to encompass all entities. For example, he characterizes the expression 'physical' (which is for him primitive) by saying that everything that is physical exists in a certain place and at a certain time. But places and times, over which he quantifies in the corresponding axioms, do not themselves appear in the diamond, and it is hard to see how they can be integrated in the uncompromising architecture of Sowa's system. They would seem to have a place *next to* the diamond, not within it. And even if physics has not yet encompassed space and time

in a Grand Unified Theory, it is indispensable for the ontologist to capture such important categories in his system.

9. *Conclusion*

Our criticisms of OpenCyc and the Sowa Diamond show that the suggestions proffered within the fields of informatics and knowledge representation for the formation of a top-level ontology are not always satisfactory. In drawing on Aristotle's list of categories, in this chapter I have developed suggestions for a top-level ontology that corresponds to the basic characteristics of Basic Formal Ontology (BFO). The three ontological dichotomies of dependent versus independent, continuant versus occurrent, and universal versus particular, form an armory of categories that, by means of further distinctions, can be built upon and refined. In fact, BFO is already being used, in applications, by a number of biomedical ontology groups, many of which are members of the OBO Foundry (see Chapter 1).

Chapter 9: The Classification of Living Beings

Peter Heuer and Boris Hennig

Biomedical ontology is the study of entities in the domain of biomedicine, specifically of the general kinds and properties which these entities instantiate. Living beings are among the key entities in this domain. No ontology of the biomedical domain would be complete which does not take into account the fact that living beings are subject to division into species and genera. One reason is that this is a fact about living beings, and the best ontology is the one that is most accurate to reality. Another reason, specific to *biomedical* ontology, is that health is species-relative; hence, it is important to understand the way in which living beings are classified and divided into species. Further, living beings are composed of parts, such as organs, many of which have specific functions. Insofar as biomedical ontology aids in the practice of medicine and clinical research, it is crucial to know which specific function to attribute to which part. But a part of a living being can be said to function or malfunction only against a background of knowledge about the features that are characteristic of the species to which this living being belongs.

This chapter proceeds in five steps. First, we will describe and justify the structure of the traditional system of species classification. Second, we will discuss three formal principles governing the development of taxonomies in general. It will emerge that, in addition to these formal principles, a division of living beings must meet certain empirical constraints. In the third section, we will show that the traditional division of living beings into species best meets these constraints. Fourth, we will argue that a taxonomic system based on this notion of species provides a more natural alternative to the many arbitrary classifications that are possible. Hence, the traditional classificatory system is also the most natural one. Finally, we will discuss and reject an alternative account that suggests defining species solely with a view to their evolutionary history. We will argue that taxonomic trees do not depict hereditary connections but, rather, something else.

1. *The Structure of the Traditional System*

The purpose of a taxonomic system is to systematize the names of, and our knowledge about, kinds of entities. In the taxonomy that is in use in present-day biology, the European domestic cat is classified as follows:

Domain	<i>Eukaryota</i>
Kingdom	<i>Animalia</i>
Subkingdom	<i>Bilateria</i>
Phylum	<i>Chordata</i>
Class	<i>Mammalia</i>
Legion	<i>Cladotheria</i>
Cohort	<i>Placentalia</i>
Order	<i>Carnivora</i>
Family	<i>Felidae</i>
Genus	<i>Felis</i>
Species	<i>Felis sylvestris</i>
Subspecies	<i>Felis sylvestris sylvestris</i> ⁴³

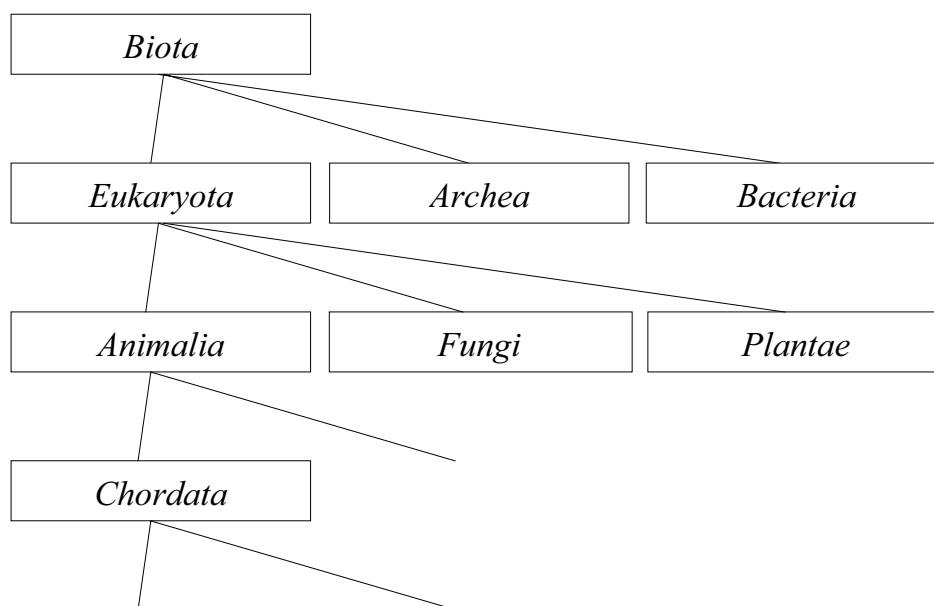
Felis sylvestris sylvestris is located at the bottom level of a series of distinctions. The domain *Eukaryota* is distinguished from other domains such as *Bacteria* and *Archaea*; the kingdom *Animalia* is distinguished from other kingdoms such as *Plantae* and *Fungi*; and so forth. As a whole, these distinctions constitute a tree-like structure; that is, a structure with one top-level node that divides into several child branches, which in turn divide into further branches. The branches at the bottom of the tree, which do not divide into further branches, are called *leaves*. The initial segment of the series of distinctions by which *Felis sylvestris sylvestris* is classified may be depicted as in Figure 1. The nodes here are called *taxa* (singular: *taxon*). All taxa above the level of Species are called *higher taxa*. The purpose of situating individual species into such a tree can be explained best by considering both technical and empirical constraints.

First, tree structures can be browsed much more efficiently than lists of items. For instance, suppose that the question is ‘To what species does a given insect belong?’ and this is to be decided by matching that insect’s features against a complete list of species descriptions. Since the class *Insecta* includes more than 750,000 known species, this will take a long

⁴³ See *Taxonomicon*. We leave out several intermediate taxa such as infrakingdom, branch etc.

time; in the worst case, it will involve 750,000 steps. By contrast, in the worst case scenario searching a tree structure with two branches at each level and 750,000 leaf nodes would only take approximately 20 steps.

Figure 1: *Fragment of a Taxonomic Tree*



Taxonomic trees have the further advantage that information associated with their leaf nodes can be stored and retrieved very efficiently. For instance, a knowledge representation that contains information pertaining to all chordata alike, at the lowest level, would contain much repeated information. It would tell us that cats have a spine, dogs have a spine, horses have a spine, and so on. It is clearly more efficient to associate information that holds true of all these species alike with a higher level node, namely *Chordata*, so that it needs to be stored only in one place. (Compare Aristotle, *Parts of Animals* 639a15–30.) Such information, then, would be inherited by the nodes lower down the tree. In this way, one may gather all information about cats by traversing the tree upwards and adding more and more general knowledge about mammals, vertebrates, chordates, etc. As a result, one may conclude from the fact that a given living being is a cat, and cats are chordates, that this living being has a spine. This only works, however, in systems where no taxon has more than one immediate parent group (as we see in Chapter 8); that is, it works only in tree structures.

It is also easy to see why a taxonomic system must have highest and lowest elements in order to be effectively browsed; that is, a root node and leaf nodes. First, a search within a taxonomic tree must proceed in only one direction, and it must begin from a point from which all taxa may be reached. Therefore, there must be a highest taxon; otherwise, the search could not begin. If there were two highest-level genera, there would also have to be a procedure to decide where to start looking, and this would amount to the introduction of one highest order taxon. Further, any process of traversing the tree downwards must be guaranteed to terminate at some point. A system that endlessly divides every taxon into further subtaxa would not be of much use.

It should be clear that a system for classifying biological kinds or species can only be established on the basis of knowledge about the *particular* living beings which instantiate them. In order to locate a species in the taxonomic tree, one must already possess extensive knowledge about the features of its instances. Further, many of the terms that are used in a scientific taxonomy are also used in everyday life and, hence, have a meaning of their own. The mere labels already encode empirical knowledge. It is not always possible to introduce higher-level taxa by focusing on one feature, as in the case of *Mammalia* (which are defined, through the presence of mammary glands in females). For instance, the division of living beings into plants and animals is used in everyday life and, hence, has a meaning of its own; but there is no single feature that all plants share and all animals lack. We will show that the traditional classificatory system is also the most *natural* one.

It is a remarkable fact about living beings that they admit of a classification in a tree structure. There may be cases that are difficult to accommodate, but it is still possible to amend the tree structure to make it fit. For instance, the platypus has features of typical mammals but also lays eggs. In order to fit it into the system, the class *Mammalia* was divided into the subclasses *Prototheria*, *Metatheria* (including marsupials), and *Eutheria* (higher mammals). The platypus is classified under *Prototheria*. This is only one instance in which an empirical discovery has led to a change in the taxonomic system.

2. Three Regulative Principles

The last section emphasized the practical advantages of taxonomic trees. Following Kant, we will now consider three formal principles governing

the construction of any such tree. It will emerge that, in addition to these logical constraints, a good taxonomy of living beings must be based on an empirically founded basic division which will turn out to be the division of living beings into biological species.

Kant postulates what he calls three formal principles of reason, which are necessary for systematizing any domain, namely, (1) the principle of specification (Kant, 1781, B 682), (2) the principle of unity (B 680), and (3) the principle of homogeneity (B 685). These principles are merely regulative (B 672). This means that they are not aspects of reality, but only guide our inquiry into the nature of real objects. Since they direct the acquisition of knowledge and are not derived from this knowledge, they may be called *a priori*.⁴⁴

The *principle of specification* demands that, for every taxon, one should ask whether it may be divided into further subtaxa. Since the same question is to be raised concerning the subtaxa, the process of division does not come to a natural end. Every species may, and should, be divided into subspecies, and these subspecies should be further divided. As a consequence, Kant claims that there is no lowest species. In his *Logic*, Kant writes:

Even though we may have a notion that we apply immediately to individuals, there may still be specific differences regarding this notion, which we either do not notice, or neglect. It is only comparatively, as a matter of convenience, that there are lowest level notions, which receive their meaning as it were by convention, as it were, when one agrees not to proceed further down. (*Logik* §11, *Akademie-Ausgabe* vol. IX, p. 97; our translation)

Note that, although the division of taxa into further subtaxa can go on indefinitely, it will never reach the level of individuals. It is possible to distinguish species in such a fine-grained manner that every individual is taken to be an instance of its own kind, but even then, the individuals will not coincide with the species to which they belong. Kant writes, metaphorically, that ‘the logical horizon consists of smaller horizons (subspecies), but not of points (individuals), which possess no extent’ (B 686). Just as we can always further divide a geometrical line without reaching the level of geometrical points, so too, we can always divide a taxon into further subtaxa without reaching the level of individuals.

⁴⁴ For a discussion of the role of *a priori* knowledge in formal ontology, see Chapter 2, Section 4.

The converse of the principle of specification is the *principle of unity*, which Kant also calls the principle of genera. It requires that we always try to bring different taxa under a common higher-level taxon, and ask further under what higher-level taxon this latter taxon may be brought. Kant writes: ‘There is a genus that cannot be a species [i.e. a subtaxon], but there is no species that could not be a genus [i.e. have subtaxa]’ (*Logik* §11, 97). The principle of unity instructs us to always proceed to higher and more general taxa, and again, it does not seem to tell us where to stop. As a matter of fact, however, there must be some point where the application of this principle comes to a halt.

To these two principles, Kant adds a third, the *principle of homogeneity*. This principle demands that, whenever we draw a clear distinction between species, we should be aware of the possibility of borderline, intermediate, or mixed cases. We should always keep in mind that, as a rule, the universe is continuous: between any two distinct entities, there can be an intermediate one (B 687). The principle of homogeneity counteracts the principle of specification by postulating a certain affinity between instances of different species. It does this in two ways.

First, for every two taxa, some common higher-level taxon can be found, however remote. Hence, the instances of every pair of different species are also alike in some sense. This follows from the principle of unity. Second, for all taxonomic divisions, there may be intermediate stages or forms. Following the 18th century biologist, Charles Bonnet, this may be dubbed the principle of continuity (Bonnet, 1766).

The principle of continuity, itself, can be understood in two different ways. First, it may be taken to postulate that, between every taxon and its higher-level taxon, further taxa may be introduced. An application of this principle has already been mentioned: the insertion of subclasses of *Mammalia* in order to accommodate for the features of *Platypus*. This procedure is quite common; for instance, in the complete classification of *Felis sylvestris sylvestris*, three taxa have been inserted between cohort and order – magnorder, superorder, and grandorder – and more still could be inserted.

A second way of understanding the principle of continuity is that, for every division of taxa into separate subtaxa, there will be certain items that fall between the cracks. In general, wherever we draw distinctions, there may be borderline cases. However, the existence of borderline cases does not prove that there are no distinctions to be drawn. Further, it will become

apparent that there are real distinctions in nature between different biological species.

Kant emphasizes that all three principles are only of heuristic value. They direct our empirical research in that they tell us where to look for further evidence. They do not directly apply to the world we experience; that is, they do not tell us that there must in fact always be a common genus, a further species, or a borderline case. The principles only encourage us always to look and see whether there are common genera, further species, or borderline cases. When Kant writes that there can be no lowest species he can only mean that there can be no *logical* reason to stop subdividing a taxon. In principle, it is always possible to insert an intermediate taxon between a given taxon and its subtaxa. However, Kant does not tell us when to stop looking. Further, since all biological individuals possess their own unique features and, since as long as there are living beings new living beings may be born, there is no limit to the possibility of dividing species into further subspecies, and no limit to the possibility of borderline cases.

Of course, there are practical and theoretical reasons why we should stop adhering to the Kantian principles after a certain point. Consider the principle of unity. It demands that we should try to bring every taxon under a higher-level taxon. This process must come to a halt, at least when the highest possible genus, ‘being’ or ‘entity’, is reached. But there are also reasons to stop applying it well before the highest possible genus. The most general set of beings relevant to biology is the set of living beings (*Biota*). To be sure, it is possible to subsume living beings, artifacts, and other physical objects under one common header; but this is of no practical value. Moreover, it tends to blur essential differences, which is an important theoretical consideration for anyone interested in an ontology of the biological domain that is accurate to biological reality.

For instance, there are no criteria of identity that apply to material things in general. Living beings remain the same entity as long as they stay alive, and they need to exchange matter in order to do so. By contrast, lifeless objects may be identified, simply, in terms of their matter. Further, although (most?) artifacts are lifeless objects, an identification of artifacts in terms of their matter leads to certain problems: a ship arguably does not cease to be the same ship when all its planks are replaced.⁴⁵ Hence, living beings, artifacts, and other physical objects should be distinguished, not in

⁴⁵ This is known as the ‘Ship of Theseus Problem’ (Hobbes, *De Corpore* XI). See, for instance, Rea, 1995.

terms of specific differences regarding their features and qualities, but in terms of the principles according to which they may reasonably be identified as the same things over a certain period of time (Scharck, 2005). This means that an ontology of the biological domain does not have much use for a common genus that embraces these different kinds of beings, although such a genus is required for an upper-level ontology such as BFO (see Grenon, *et al.*, 2004; Grenon and Smith, 2004; Grenon, 2003).

We conclude that the three principles put forward by Kant apply to all taxonomic systems but that, in each case, they need to be complemented by empirical constraints. In biology, there is a particular highest taxon (*Biota*) and a basic level of division on which the whole taxonomy of living beings is founded. It is important to keep logical and empirical constraints distinct from one another. It is an empirical fact that all living beings have something in common, so that they constitute a realm that admits of a taxonomic classification. It is also an empirical fact that there is a point where the division of taxa of living beings into further subtaxa comes to a natural end. That such a basic division exists is not a logical requirement. It is a logical requirement, however, that taxa divide into further subtaxa such that a tree structure results.

In accordance with the three principles of classification named by Kant, one may establish this structure by proceeding both upwards and downwards: upwards by grouping species together in higher taxa and by bringing the higher taxa under taxa that are still more general; and downwards by dividing the realm of living beings into domains and subtaxa. The most general distinction we make within the realm of biology is the one among *Bacteria*, *Archaea*, and *Eukaryota*. From this point on, one may develop the system by introducing a series of distinctions. At the same time, however, the system is supposed to capture the known species of living beings. To this end, one should look at the accounts and descriptions of different biological species such as *yarrow*, *cat*, and *sparrow*, and consider how they are best grouped together under more general labels. The task is to unify and merge different groups into higher order groups.

As has already been noted, the advantage of this bottom-up procedure is that we may associate certain bits of knowledge with the higher order groups instead of redundantly associating them with several lower level groups alike. This procedure facilitates the learning and teaching of facts about kinds of living beings. The purpose of a classification of living beings is to provide a basis for the storage and acquisition of knowledge

about living beings, not to merely impose order. This is even more important in contexts where knowledge is processed automatically, and where vast amounts of knowledge are maintained. It can cause a great deal of trouble to maintain and update a system containing redundant information. For instance when new information concerning all insects comes to light, data would need to be changed in almost a million different places in the same way. But if the information is stored only in one place, namely under the label *Insecta*, such a change is easily made in one step.

Though the classification of kinds into higher order kinds has such a practical purpose, it must not be arbitrary. Indeed, the best classification will always be one that refers to features that are, in fact, typical for the respective range of living beings. The class of mammals is a group of items that belong together in more than one respect, whereas the introduction of a class of two-legged animals would soon cause trouble (since it would include birds and humans alike). Which divisions are appropriate can only be seen by simultaneously pursuing the downward movement of division and the upward movement of unification.

3. *Biological Species*

We will now argue that a system for classifying living beings must be based on a division into biological species. This gives rise to the question of what a biological species is. This section will provide an answer to this question.

In a logical sense, every group that may be divided into subgroups is a genus, and every group that may be brought under a higher order group is a species (Kant uses the terms in this sense in the second passage quoted above). The biologist, however, uses ‘species’ in a much narrower sense. Biological species constitute only one level within biological taxonomy.

Below the level of biological species, one may distinguish populations, varieties, races, and forms; but these distinctions are always, to some extent, arbitrary in that they involve merely geographical and phenomenal differences. The taxa above the species level differ from species in that they are only associated with a fragmentary or ambiguous description. There are instances that satisfy all and only the criteria that apply to the species *Felis Sylvestris*, but there are no instances that would satisfy only the criteria that apply to the class *Mammalia* in general. *Mammalia* is an abstract taxon. There are different kinds of mammals; some have fur, long tails, exposed genitals, and some do not. The description of the class

Mammalia is incomplete in this regard. Every mammal is necessarily an instance of some species, whose description can be made complete to an arbitrarily detailed degree. This does not mean that the class *Mammalia* does not really exist, let alone that there are no mammals. It means, however, that there are no mammals over and above the instances of particular species of mammals.

In this respect, class names are like mass terms. Mass terms such as ‘milk’ apply to real things, but they do not refer to countable items. In reality, however, everything that exists can also be counted: every instance of milk is an instance of so and so many centiliters of milk, and centiliters of milk can be counted. Nonetheless, it makes perfect sense to speak of *milk*, in contrast to definite portions of milk. When we do so, we abstract from the countability (portioned nature) of all real milk. Likewise, class terms such as ‘mammal’ apply to real things, although in fact, every mammal is also an instance of some more specific species. When we use such terms, we abstract from certain specific features of a living being.

Species provide the units of biological reality, and taxa below and above the species level can only be introduced against the background of a species division. Therefore, it is of the utmost importance to be clear about the precise circumstances under which a taxon constitutes a species. That the discovery of a new species is something biologists tend to be proud of shows this is important as well. In some cases, species bear the name of their discoverer, as for instance the *Ophrys regis fernandii* or the *Epipactis mülleri*, named after their respective discoverers, King Ferdinand and Müller. To discover a new biological species is regarded as a lasting achievement. However, since it is always possible to introduce further distinctions, it is logically possible to divide every known group of living beings into further subgroups. The question is, under what conditions is such a division, in fact, a division into different biological species, rather than a division into arbitrary sets, higher taxa, or parts of the same species. For instance, many plants differ from others merely because of the quality of the soil, or only because of their geographic location; but such differences should not license a species distinction. There should be a limit to making divisions since, after all, biological species have to be registered, described, learned, and taught. In order to avoid proliferation of species divisions, one needs non-arbitrary and ontologically sound criteria for what biological species are. Ideally, what we need is a basic division of living beings into species that carves reality at its joints (Plato, *Phaedrus* 265e).

Yet the question of what biological species are is subject to considerable dispute.⁴⁶

A perusal of the history of philosophy and science shows that the notion of a biological species was uncontroversial until, roughly, the 1850s. With the advent of Darwin's theory of evolution, according to which forms of life are subject to constant change, the claim that biological species are part of a natural order becomes problematic. All clear distinctions between species seem to be temporary, and the criteria according to which they are drawn begin to appear arbitrary. It is no wonder that, as a consequence, there are divergent opinions as to what counts as a biological species and a good classificatory system.

As we will see, however, a closer look reveals that many of the different accounts of what biological species are, in fact, do not contradict each other. They are not all of equal importance, and they are systematically related in such a way as to complement one another. In order to determine what biological species are, we need to consider two things. First, living beings maintain and reproduce themselves. Therefore, it is quite natural to assume that a species is a group of individuals that is engaged in generating and breeding further members of this group. The idea that species are basically reproductive communities has been put forward by Ernst Mayr (Mayr, 1996).

Second, reproduction and self-maintenance can be successful or not and, where they occur, there must be certain standards according to which their success may be measured. When one spells out these standards, one ends up with a description of a prototypical and idealized (canonical) instance of the respective species. A cat reproduces successfully if the result of what it does is something that satisfies all criteria that apply to healthy and typical cats. This motivates the account of biological species suggested by Plato and Aristotle. Species are associated with standards of typicality, and to describe a species is really to describe its ideal case: the idea (*eidos*) or essence of its instances.

One can bring together both strands in the following characterization: *Biological species are universals instantiated by members of reproductive communities that secure the (at least relative) permanence of a form of life that is characteristic of members of this community, by passing it on to their offspring.*

⁴⁶ See e.g. Ghiselin, 1974; Hull, 1997; Mayden, 1997; Ereshefky, 2002; Reydon, 2005.

Species are not sets of living beings; therefore, some biologists like Ghiselin (1974), Caplan (1981), and Hull (1997) have claimed that species must be individuals. This claim, however, rests on the assumption that the alternative between sets and individuals is exhaustive, which holds true only in an ontology such as the one suggested by Quine. Quine's ontology, however, is suitable for mathematical and physical entities only, and not for living beings. Species are neither individuals nor sets, but universals.

More specifically, to instantiate a species is not only to exemplify a set of characteristic features, but also to lead a certain life. For example, the instances of *Felis sylvestris* are born in a certain way, develop in a certain way, and perform certain characteristic activities during their lives. What they typically do in the course of their lives does not only contribute to their life; rather it *constitutes* their typical life. A description of what is characteristic of cats cannot consist in mere a list of features, but only in a story about the typical life of a cat (Thompson, 1995).

It is important to note that individual instances of a species may transmit their characteristic features to their offspring even if, for some contingent reason, they do not possess them. For instance, a cat with three legs, in most cases, will generate offspring that has four legs. In reproducing, instances of biological species do not just copy their own particular makeup, but transmit a form of life that is characteristic of instances of their species. Therefore, a species is constituted by all individuals that may successfully reproduce, such that instances of the same form of life result.

That species are instantiated by reproductive communities does not imply that all instances of a species can actually mate with all other instances of this species. First of all, it is not necessary that all instances of a species do, in fact, successfully mate with all other instances. Two male individuals of the same species cannot mate and generate offspring, but they both can in principle generate offspring by mating a female instance of the same species. Second, individual instances of a species may be entirely infertile, raising the question of whether they belong to the same reproductive community. But all that follows from our understanding of species is that, for all instances of a species, not to be able to generate further individuals with certain characteristic features constitutes a *defect*. If an individual is infertile, it thereby fails to belong to the species only if its infertility does not constitute a defect; and whether infertility is normal or pathological can usually be ascertained by independent means. It is also a matter of dispute whether two populations that cannot interbreed because

of geographical barriers constitute a species or not. In such a case, it is not clear whether both populations actually belong to the same species until it can be shown whether, in principle, they are able to interbreed.

These details do not alter the general idea that species are instantiated by reproductive communities of individuals. In order to flesh out this idea, one may describe what conditions must be fulfilled for individuals to successfully reproduce and preserve their characteristic form in more detail. This can be done by further discussing how a population manages to ward off distorting influences, and how reproduction works; for instance, by providing a detailed account of how genetic codes are merged and copied. Knowledge about genetic processes may be adduced in order to explain how living beings actually manage to transmit a characteristic form of life to their offspring. Such an explanation of how reproduction works complements the account developed so far; it does not lead to a different account of what species are.

However, the suggestion to define species *merely* in terms of evolution is problematic in certain respects. Species of higher forms of life are not rigid but, instead, provide for a certain range of differences concerning the features, form of life, and behavior of their instances. Thereby, they also allow for the development of new features that may be distinctive of certain races, forms, or varieties. But the emergence of a race should be distinguished from the development of a new species. Races are only possible within the range that is left open by the proper description of a biological species. For instance, the proper description of *Felis sylvestris* leaves open whether its instances have black or white fur. The coming into being of races, forms, and varieties is not an instance of evolution but, rather, the realization of features or forms of life that instances of some already existing species can exhibit. Races may remain stable for contingent reasons, but they tend to disappear when their instances interbreed with other instances of the same species.

The development of races can explain the emergence of new species only if additional conditions hold; for instance, that the members of a race have been isolated and have changed because of inbreeding. Long isolation might lead to a radical change in reproductive behavior, so that interbreeding with other instances of the same species ceases to be possible. Such isolation, however, should be taken to abolish the unity of the original reproductive community and, where this unity is compromised, the permanence of the form of life characteristic for a species is not granted. As a matter of fact, species need the possibility of crossbreeding

between as many different populations and individuals as possible in order to retain their form of life. When a significant portion of a reproductive community ceases to contribute to the reproduction of the whole species and begins to constitute its own species, both parts of the original reproductive community come to instantiate a *new* species (although one of them may retain the old name), having evolved from the old one.

But this does not mean that species are changing things. They are universals. When something turns from red to green, the universal that it exemplifies does not change; that is, Red does not turn into Green. Rather, the thing changes by coming to exemplify another universal. Likewise, when a population comes to instantiate a new species, it is not the species itself that changes, but the population that ceases to instantiate one species and comes to instantiate another. Of course, we can say that a species changes in the same sense in which we can say that the color of an item changes when it turns from red to green. But this does not mean that the species itself undergoes a change, just as the change of color is not a change that the color Red undergoes.

Further, the process by which a population may come to instantiate a new species cannot be a continuous one. First, a *continual* evolutionary development could only take place where the evolving beings do not divide into biological species at all since, during this development, genetic changes are transmitted to the offspring without correction. There would be no difference between successful and failed reproduction and, hence, there would be no form of life that would be characteristic of the living beings in question. Second, even where evolutionary change does not occur continually but only temporarily, the criteria of successful reproduction are suspended as long as the change is taking place. As long as a species evolves, no one could possibly tell whether its offspring is as it should be; since by assumption, this offspring exhibits a new form of life, and this new form of life might become characteristic of resulting populations.

This implies that there can be no purely evolutionary concept of a biological species.⁴⁷ Where there are species, there is no evolution, and where evolution takes place, there are no species. A species can be the result of evolution and the starting point of more evolution, but as long as evolution is taking place, there are no clear differences between features that are characteristic for the evolving beings and features that are not, and hence there is no species.

⁴⁷ Pace Hennig, 1966; Kornet and McAllister, 2005; Griffiths, 1996; and Millikan, 1999.

The relevance of these considerations becomes obvious when we consider that the permanence of a biological species is the conceptual precondition for a taxonomic system such as the Linnaean one. The Linnaean taxonomy systematizes universals, not populations; and whereas populations can change with respect to the universals they instantiate, the universals themselves do not evolve. On the other hand, the existence of continuous change is one of the central assumptions of Darwinian evolutionary theory. Hence, there are conceptual reasons why Linnaeus denied the possibility of evolution, and why Darwinians, on the other hand, have problems with the concept of a biological species. This conflict, however, is only apparent. Evolutionary theory does not really describe how species undergo a change; it only describes how populations come to instantiate new species.

This does not at all diminish the importance of evolutionary theories to taxonomy. In particular, it does not mean that evolution could not explain *why* and *how* living beings divide into biological species. It only means that evolutionary theory cannot provide the whole and exclusive basis for a taxonomic division of living beings into species.

It should be clear that we need, at least, the concept of a *relatively* permanent species in order to do taxonomy. The process of dividing taxa into further subtaxa can only be brought to a halt if we assume that there are biological species with *certain* stable characteristics. We can do so by admitting that species may change, but abstracting from this fact and only considering the results of these possible changes at one instant of time. In fact, this is all we need since we are only interested in a classification of the living beings and the results of evolutionary change at a certain instant of time.

4. *The Search for a Natural System*

A system is artificial if it distinguishes between different kinds of things according to criteria that provide a superficial overview of the various forms of life, in reflection of chosen purposes. In order to establish a natural system, we need to inquire into the natural and objective order of things, so that we may divide our domain by criteria that are founded on the nature of the things to be ordered and, thereby, provide a better alternative to the many arbitrary classifications that are possible.

To be sure, in several contexts it is useful to classify living beings according to criteria that refer to our own purposes. Such classifications are

already found in the Old Testament, where animals are distinguished into pure and impure ones, and members of the religious community can only eat the pure ones. One may also classify animals according to where they may be hunted, and what side dishes or kinds of wine fit with them. Such classifications, however, are valid only relative to certain human communities; they refer to things that are of more or less value to the members of specific groups of humans. Where a unified and scientific classification is in order, it does not make much sense to choose criteria that may vary from one set of scientists to the next.

One might object that a natural system is not needed, since it would be enough if scientists agreed to use some fixed set of arbitrary criteria. These criteria should not vary; but they need not be natural. Why do we need a natural system for classifying living beings? The answer is that an arbitrary set of criteria may become obsolete for irrelevant reasons. If the agreement of all scientists to use given criteria is itself arbitrary and not founded on objective facts, all scientists might as well decide to change the criteria for arbitrary and contingent reasons. Natural systems can only fail for relevant reasons, that is, only when reality changes, or if they were inadequate (that is, not truly natural) in the first place. Moreover, it is unlikely that all biologists would agree on a common set of arbitrary criteria, since different biologists (botanists, geneticists, physiologists, etc.) pursue different projects and take different views on biological reality. In fact, there have been a wide variety of different and even incompatible classificatory systems in biology before a natural system was established.

Further, it may be objected that every system of classification, including the biological one, is in some sense artificial. After all, science is an artifact, and so is every scientific taxonomy. There is some truth to this objection. Science is done and maintained by humans; however, this does not mean that the results of science are arbitrary. The traditional, Linnaean, biological taxonomy is based on a division into biological species that is found in nature, and is constrained by empirical facts. The task is to find out what is really essential to specific forms of life, and how different species actually differ from and are similar with one another.

Finally, one might object that the criterion of cross-fertility is as arbitrary as any pragmatic criterion that is used by scientists in order to suitably systematize their domain. After all, humans are especially interested in breeding plants and animals, and this may be why cross-fertility is so important for them. It may also seem that, as Kant says, scientists assume the existence of species only as a matter of convenience,

in order not to be forced to constantly divide all their taxa into further subtaxa. This view, however, would overlook the important fact that the instances of species themselves ensure the permanence of a certain form of life. Reproductive communities are engaged in generating offspring with certain characteristic features and, in this sense, they are engaged in perpetuating and stabilizing their own species. Put differently, there is an objective division of living beings into species because there are objective limits to reproduction; and there are species to be distinguished because reproduction occurs within these objective limits. Without limits to reproduction, instances of different species could mix and generate indefinitely many intermediate forms. In this case, it would be difficult to tell whether reproduction is successful or not. This, however, is not the case.

This is why the criterion of cross-fertility is more powerful than other criteria by which we distinguish kinds of things. It yields divisions that are probably only superseded in their clearness by the distinctions we draw between different individual objects. Individual objects may be distinguished from one another as long as they occupy a clearly limited location in space and are impenetrable in some sense, so that they do not merge with other objects, and do not move discontinuously. Similarly, particular species may be distinguished from other species (1) because their instances do not successfully interbreed with instances of other species, such that the boundaries between different species are impenetrable, and (2) because all instances of a species derive from ancestors that belong to the same species, such that there is a continuous path that leads from one instance to the next one. Most importantly, the existence of reproductive communities implies that the realm of living beings is not a continuum. There are real distinctions between different species because there are real reproductive barriers. Kant's principle of homogeneity does not apply.

The biological classificatory system is not natural in the sense that it may, as a system, be found in nature. It is natural because and insofar as humans have established it according to objective criteria that reflect the nature of things, and not according to arbitrary and artificial ones.

5. Taxonomy and Ancestry

In biology, attempts are being made to define species and higher taxa only by reference to the common ancestry of their elements. Some biologists

have suggested that the evolutionary tree of descent directly mirrors the division of living beings into species. This has led to the idea that the tree of ancestors of biological species is also the best system for classifying them, and that it is evolutionary biologists and paleontologists, rather than taxonomists, who should lead the search for a natural system (see Mayr, 1969). But evolutionary theorists and paleontologists are concerned with establishing a family tree of biological species and, as we will argue, family trees and taxonomic systems are fundamentally distinct things.

Evolutionary biologists claim that the traditional classificatory system is not a natural one. We have already seen that for a classificatory system to be natural, it needs to be made according to non-arbitrary criteria which match the nature of things. A phylogenetic tree may seem to be more natural, in this sense. The question is thus whether, from a logical point of view, it makes sense to replace the traditional classificatory system with a new one based only on common ancestry.

This question already presupposes that we are able to give a reasonably complete family tree of biological species. Since such a tree cannot be based on direct observation of presently existing forms of life, the main method for establishing such a tree is a comparison between extant forms. However, similarities between living beings of different kinds, at best, indicate that they might have common ancestors. A method that allegedly serves to discover hereditary bonds was developed in the 1950s by Willi Hennig, who aimed at establishing a *cladistics*; that is, a classificatory system that is exclusively based on phylogenetic kinship (Hennig, 1966). To this end, particular features are singled out by way of comparison, and used in order to establish so-called *cladograms*. The comparison is carried out on the basis of morphological features, characteristics of the digestive system, and the DNA sequences of extant species.

In order to establish cladograms, derived features are distinguished from non-derived ones. Non-derived features are supposedly older in terms of evolution; they are also features of the ancestors of the species under consideration. For instance, it is a non-derived feature of mammals that they possess a spine, since instances of other and evolutionarily older classes also have spines. The derived features of a species, in contrast, are assumed to be younger in terms of evolution and occur only in this species as we find it today. The totality of derived features constitutes the *principal form*, which is considered a possible candidate for an evolutionarily older and more original species. A derived feature of mammals is that they possess a placenta, and the assumption is that all mammals derive from a

species that was marked by the possession of a placenta, among other features. Species that agree with respect to their derived features are taken to be cognate. In this way, one can establish trans-specific types similar to those suggested by 19th century biologist Georges Cuvier (Cuvier, 1827, *Introduction*). The assumption that the classification that is established by such means no longer rests on morphological criteria is, however, purely hypothetical.

Further, even if we were in possession of an adequate, complete, and empirically founded family tree of biological species, this tree would not depict the system of biological kinds. The reasons for this are conceptual ones. Genera, families, and other higher taxa cannot be ancestors of the extant forms at the same time, since the extant and the ancestry forms are all biological species. For instance, the *Archaeopteryx* is probably the ancestor of all kinds of birds known today, but *Archaeopteryx* is a species and not a genus, family, or class. In a family tree, the *Archaeopteryx* would be represented by a node whose child nodes include all extant kinds of birds. In a classificatory tree, the node representing the class *Aves* would occupy this position. It should be clear that the species *Archaeopteryx* cannot be identical to the class *Aves*, for it is also a species falling under this class. Just as an individual living being cannot at the same time be its own species, an individual species cannot at the same time be its own genus, family, or class. This is so even if a taxon contains exactly one species, since an individual that is the only instance of its kind is not thereby identical to its kind. This distinction between species and their instances, and classes and the respective subtaxa, may be less obvious when the taxonomic tree is read in set theoretic terms. In cases where a class only contains one species, the set of instances of the class is identical to the set of instances of the species. But this is *not* how one should understand the taxonomic tree (Buck and Hull, 1966).

Regarding morphological similarity, the *Archaeopteryx* is especially unsuitable as a primordial or paradigmatic form, because it lacks essential features of birds. Many of the generic statements about birds do not apply to the *Archaeopteryx*. For instance, it does not yet have the large sternum that is typical for all extant birds. Hennig seems to be aware of this problem, since he explicitly neglects fossils and only compares extant species to each other.

We conclude that classes and species are related *conceptually*, rather than by way of ancestry. Evolutionary trees depict the historical sequence of generations of individual living beings; that is, the hereditary lines.

Taxonomies bring living beings under general concepts according to their features, and their purpose is to provide an order that is as clear as possible, in order to systematize biological knowledge, so that certain propositions can be inferred in the way that has been described by Cuvier.

6. *Conclusion*

We have pointed out that an understanding of biological taxonomy is essential to biomedical ontology. The most appropriate account of the division of living beings into kinds, we have argued, is provided by the traditional, Linnaean system. First, the traditional classificatory system satisfies important logical and empirical conditions for any such system. It constitutes a tree, and can therefore be quickly and efficiently browsed by both humans and machines. Further, it embraces all known species and, thereby, provides a structure for systematizing and encoding our knowledge of all biological species. Finally, it serves to determine the species of individual living beings effectively. A mere list of forms of life, as suggested by Bonnet, does not allow for this; it would be extremely tedious to browse.

We have further argued that, in order to establish a classificatory system of living beings, it is not enough to adhere only to the logical principles that govern all possible taxonomies. Other conditions that have to be met are that (1) the taxonomic system must be founded on a basic division, such that the division of taxa does not go on indefinitely, and (2) the classificatory divisions within the system must be reliable and non-arbitrary. A classificatory system is a candidate for a possible taxonomy of living beings only if the basic and the higher-level divisions accord with the facts.

A division of living beings into biological species provides the basis of the traditional system. This division is well founded, since it mirrors the reproductive barriers between individual living beings. The boundaries between different species do in fact exist; they are the reproductive barriers that prevent the interbreeding of individuals from different species. In biology, the highest taxon is the group of living beings. No higher taxon is needed. The classificatory divisions in between are also well-founded and non-arbitrary. For the higher taxa are intended to correspond to essential features that instances of certain species share, and by which they differ from instances of other species. In this way, a hierarchy of higher taxa is established. It is important not to restrict attention to only a few features or

bodily parts of living beings in order to classify them; instead, one should always consider the living beings as wholes, taking into account their visible makeup as well as their inner structure. Even the DNA, however important it is in modern biology, should only be considered as one feature of a living being among others, which adds to the overall picture. Living beings that belong to different classes and differ widely with respect to their phenotype often possess surprisingly similar genotypes.

Further, we emphasized that family trees and taxonomic systems are fundamentally distinct things. Taxonomies systematize living beings according to shared and distinctive features, and their aim is to provide a clear and effectively usable system for describing and identifying living beings. Higher and lower taxa are related conceptually, and not in terms of ancestry. Evolutionary trees, in contrast, depict hereditary lines among different species, just as a family tree represents the pedigree of an individual. Taxonomic and hereditary relations have a different logical status, and neither can be reduced to the other.

We conclude that, in biomedical ontology, the traditional taxonomic system as developed by Aristotle and Linnaeus remains indispensable. Hereditary trees may be of help in establishing such a system, but they cannot replace it. To be sure, facts about the ancestry of a species should always be accounted for and acknowledged in taxonomy. A system that does not group species together – when, in fact, they have a common ancestry – would not be a natural one. But this does not mean that ancestry is the only relevant criterion, or even that evolutionary theories alone can do the job. It certainly does not imply that the genus of a species coincides with its ancestry.

Chapter 10: Ontological Relations

Ulf Schwarz and Barry Smith

1. *Formal Ontological Relations: What are They? What are They For?*

During our discussion of the ontological sextet and the associated classification of reality in Chapter 8, we saw that the classical Aristotelian square can be extended to encompass both particulars and universals, and both continuants and occurrents, to yield a total of six categories. We had postulated very specific relations among the entities in these categories, such as the relations of inherence, participation, instantiation, exemplification, and characterization. In this chapter, we will discuss how to characterize the relations among the entities in these categories more precisely and to define the relational expressions used in ontologies in a more rigorous and unambiguous way.

The relations mentioned are genuine formal-ontological relations. That means, first, that they are *ontological*, not merely logical. Certainly, as thinkers such as Frege and Boole have argued, studying logic can help us to gain some understanding of general aspects of reality (see Meixner, 1992). However, to gain a more complete understanding of this reality it is ontological relations, not logical ones, which we must study. Second, it means that the given relations are *formal*. They are not additional components of reality; rather, they are that which binds existing entities into larger unities (Ceusters, *et al.*, 2006). An example of a formal-ontological relation in this sense is the *part of* relation, in which your arm, leg, head, and so forth, stand to the whole that is your body. In order to explore and understand the ontological structure of reality, it is necessary to create an assay of such important relations.

The most appropriate method used for attaining this goal is to concentrate on an ontology describing some specific domain, such as biological reality (for the moment, we will leave top-level ontologies aside). Through examples derived from a delineated area of knowledge or science, we can show which relations obtain among the entities of that domain, and how the relational expressions used in corresponding ontologies can be defined in a way that is rigorous, unambiguous, and consistent. For this purpose, we must explicitly acknowledge the existence of certain other kinds of entities, such as points in time and spatial regions, which are important for expressing simple relations in biology and thus also in ontologies of the biological domain. These entities do not appear in

the top-level categories delineated by the ontological sextet. They appear at lower, more specified, levels of categorization.

The relations that appear at the top-level categories of the ontological sextet can be distinguished from one another on the basis of whether their relata are universals or particulars. We can distinguish three typical cases:

<*universal, universal*>: Both relata are universals. An example of this type of relation is characterization, or the subsumption (*is_a*) relation which obtains between the universal *human* and the universal *mammal*, such that: *human is_a mammal*.

<**instance, universal**>: The first relatum is a particular, the second is a universal. An example of a relation of this type is the instantiation relation, which obtains between this particular person named Peter and the universal *human*, or between Peter's life and the universal *life*. Another example is the relation of *being allergic to* that exists between Peter and the universal *aspirin*.

<**instance, instance**>: Both relata are particulars. Examples include the inherence relation, or the participation relation which obtains between Peter's life and Peter, or also – independently of the ontological sextet – the part-whole relation on the level of instances, which obtains between this particular nose (Peter's nose) and this particular head (Peter's head), and between both of these and Peter.

For what follows, we want to introduce two terminological conventions. First, in order to avoid any ambiguities, we will use *italics* when referring to relations which obtain only between universals, and **bold face** to express any relations which have among their relata at least one particular. Second, we will use expressions common in applied ontology: for example, *is_a* for the subsumption relation and **instance_of** for the instantiation relation. Third, we will confine ourselves to binary relations, though the ideas here expressed can be generalized in the obvious way to relations with any number of relata. Some relations, such as the participation relation and the part-whole relation, obtain in different forms both at the level of instances and at the level of universals: thus, we can speak of **participates_in** and **part_of**, as well as of *participates_in* and *part_of*. Where needed, we shall refer to these as instance-level and type-level relations.

The relations in the ontological sextet also play an important role in an ontology of the biological domain. We will use both instance-level relations, and relations obtaining between instances and universals, in order to define expressions for those relations at the level of universals⁴⁸ that are of special interest to ontologies in the biological domain such as the Gene Ontology. It is these type-level relations that are at the center of interest for anyone wishing to use information systems to represent the knowledge that is captured in biomedical science. Such knowledge is knowledge of universals. It is not *Fury* or *Black Beauty* that are described in zoology textbooks, but rather the type *horse*: not particulars, but universals and their relations of universals one to another.

Biological ontologies were developed in order to serve as controlled vocabularies for the expression of the results of biological research. Such vocabularies contain sentences of the form *A relation B*, in which ‘*A*’ and ‘*B*’ are terms in a biological ontology and ‘*relation*’ is a placeholder for ‘*part_of*’ or a similar expression. Such sentences convey general information about the corresponding biological universals. The question of which entities⁴⁹ count as biological universals (and which do not) cannot be answered easily. For the moment, we can content ourselves with the many examples of terms – such as *cell* or *reproduction* or *axon* or *glycolysis* – in ontologies which are used in biological literature to refer to general types of objects or processes. While unavoidable in biological research, statements about particular instances of these types (for example, a statement about the specific weight of this particular organism in a particular petri dish at some particular time) do not belong to the general claims of biological science. Accordingly, relations in biological ontologies link universals with other universals.

Instances are important to ontologies, nonetheless: for we cannot define what it is for universals to stand in a specific relation (for example, the type-level relation: *part_of*) without reverting to consideration of their corresponding instances. We can specify what it means for *retina* to be *part_of eye* only when we recognize that this is a statement to the effect

⁴⁸ To avoid unnecessarily complicated linguistic constructions, we will use ‘definition of a relation’ and ‘definition of a relational expression’ synonymously, except in places where this might result in confusion.

⁴⁹ In the following, the word ‘entity’ will be used as ontological term of art to refer to everything which, in some way, exists (all continuants, processes, functions, structures, places, times, etc., at the level of instances as well as the level of universals).

that all instances of the universal *retina* stand in a **part_of** relation to certain corresponding instances of the universal *eye*. This dependence of type-level relations on relations between their corresponding instances forms the basis of our definitions below of the relational expressions that link general terms to each other. For general terms are no more than the names of universals, and science is concerned with general statements about such universals, rather than with the way the world happens to be at some time or place.

In the following, we lay out rigorous definitions of the subtype-, parthood-, participation-, and location-relation between universals. Then, we introduce a method for defining other ontological relations between entities in other domains of reality.

2. Benefits and Problems in Defining Formal Ontological Relations

The consistent use of rigorous definitions to characterize formal relations will be a major step toward enabling information scientists to achieve interoperability among ontologies in support of automated reasoning across data derived from multiple domains. For, if a fruitful exchange of information is to be possible between such ontologies and the data annotated in their terms, then each of the various systems involved must treat their relations in the same way. A relational expression must always stand for one and the same relation, even when it is used in multiple ontologies.

With regard to the most basic ontological relations (such as the subtype- or class-inclusion-relation and the part-whole relation), it is apparent in the literature that the same expressions are not always taken to stand for equivalent relations from one ontology to the next. In some places, this confusion goes so far that these basic relations are not distinguished from each other at all. For example, at one time, the medical terminology database Unified Medical Language System contained the statement *plant leaf is_a plant* and the SNOMED CT vocabulary contained the statement *both uteri is_a uterus* (see UMLS, see National Library of Medicine, 2006; *SNOMED*, 2007).⁵⁰ The methodology that we are introducing will enable us to provide rigorous and unambiguous definitions for relational expressions between general terms, so that the meaning of these expressions can be stated precisely and used consistently.

⁵⁰ By the time this chapter reaches publication, it may be that some, or all, of these statements have been rectified.

In much of the literature on knowledge representation (see for example Fellbaum, 1998), relations between universals are presented as simple or basic, without any further specification. But this sort of treatment ignores essential aspects of the way in which instances in reality relate to one another. It is often not a trivial matter to determine whether or not a certain relation, in fact, obtains between the corresponding universals. For example, because there are female as well as male humans, we can certainly assert that *human testicle part_of human*, but not *human has_part human testicle*. Because there are non-human mammals with hearts, *human has_part heart* is true, but not *heart part_of human* (a pig's heart, for example, is a part of a pig). For similar reasons, it is not true that *growth has_participant human*. The temporal dimension can also be a source of special problems in determining whether relations apply. For example, even though every instance of the universal *adult* did at one time instantiate the universal *child*, it is not true that *adult is_a child*.

Unfortunately, many ontologies in the biological domain contain expressions for relations between universals whose correct usage can only be discerned through hints and loosely formulated suggestions. Without rigorous definitions, the logical connections between relations remain in the dark, and relational expressions become subject to multiple interpretations both within the same ontology (compare Ceusters, *et al.*, 2004) and also from one ontology to the next (Smith, *et al.*, 2004).

3. *Types of Relations*

Any relation to which our methodology can be applied (including the examples we discuss) must fulfill the following four criteria.

First, the relations in question must be genuine ontological relations. This means that they obtain between entities in reality, independently of our experience or methods of learning about them. They are also independent of the ways in which we represent them or make our knowledge about them processable by computers. This is by no means the case for all relations. For example, the Gene Ontology uses the relation *A annotates B*. This relation is used to link certain genes or gene products with expressions from a controlled vocabulary. This relation does not link universals of biomedical reality as they are in themselves; rather, we use it to assert that a certain link has been effected by humans between a term in an ontology and (for example) some protein. Thus, it is not an ontological relation.

Second, the relations of interest here are those domain-neutral relations which could appear, in principle at least, in any biomedical ontology. Certainly there are relations which hold between entities of specific types belonging to some specific domains. An example of this sort of specific relation would be *A is_genome_of B*, which might be used in a gene-sequence ontology. Our strategy, however, is to define, as far as possible, a small set of high-level relations in a domain neutral way, and to construct definitions for low-level domain-specific relations on this basis.

Third, the relations must obtain universally. A statement of the form *A relation B* must obtain for all instances of *A*, and not just (for example) for some statistically representative selection. In many cases, the relations with which we have to deal will express analytic connections between universals, which is to say connections that can be understood to obtain universally when we analyze the corresponding general terms with which these universals are represented: for example, in cases such as *skin cell is_a cell* or *heart attack has_participant heart*. Propositions expressing this kind of connection are true solely in virtue of the meanings of the terms involved and of the expressions which connect them. We do not have to examine the world in order to find out that, if something is a skin cell, then it is also a cell. While human beings do not need to be instructed on such matters, such instruction is needed by the automatic reasoning systems towards which ontologies are addressed.

Fourth, the relation must be definable in a simple, yet rigorous, way. This criterion is important since there are many relations that ontologists use for which they have, at best, only intuitive definitions or, sometimes, no definitions at all. Consider the relations *physically_related_to* and *functionally_related_to*, from the UMLS Semantic Network (National Library of Medicine, 2006). The former is defined as meaning: ‘related by virtue of some physical attribute or characteristic’, while the latter means ‘related by the carrying out of some function or activity’. In neither case could we apply these definitions effectively in such a way that they would help in determining whether a given example is or is not an example of one or other of the relations in question. In neither case could we use the resultant assertions for computer-aided reasoning.

We insist upon this fourth criterion because we want to introduce definitions that are easy to understand, and which can be used effectively by humans. But, at the same time, we need definitions that can be used to support logic-based computer reasoning processes and help to support consistency from one ontology to the next.

As we mentioned above, a biological ontology should represent only type-level relations of the form $\langle \textit{universal}, \textit{universal} \rangle$. The creation of rigorous definitions of these relations requires, however, first an understanding of those more basic relations which are of the forms $\langle \textit{instance}, \textit{instance} \rangle$ or $\langle \textit{instance}, \textit{universal} \rangle$. A $\langle \textit{universal}, \textit{universal} \rangle$ relation obtains only because a certain $\langle \textit{instance}, \textit{instance} \rangle$ relation obtains between the universals' respective instances, or because a certain $\langle \textit{instance}, \textit{universal} \rangle$ relation obtains between a universal and its instance. This sort of characterization of the relations between universals is anchored in the Aristotelian interpretation of the notion of universal, according to which universals exist only in their instances.

4. *Types of Relations and Limitations to the Use of Relational Expressions*

Before we can begin to define some basic relational expressions for a biological ontology, we must specify which expressions we will use to designate entities that stand as relata in corresponding relations. To do this, we must be able to speak of instances, as well as universals, in an appropriate way. We will avail ourselves of the tools of logic, including variables and quantifiers (see for example Hodges, 2001, and Chapter 5 above). Variables of various sorts are placeholders, respectively, for instances and universals of continuants, processes, and points in time:

C, C_1, \dots, C_n	stand for continuant universals;
P, P_1, \dots, P_n	stand for process universals;
c, c_1, \dots, c_n	stand for instances of continuants;
p, p_1, \dots, p_n	stand for instances of processes;
r, r, \dots, r_n	stand for three-dimensional spatial regions;
t, t_1, \dots, t_n	stand for points in time.

Continuants and processes form mutually exclusive categories. Continuants can be material entities such as a molecule, a cell, or a human being; but they can also be immaterial, such as a hole or a conduit. Immaterial continuants have some traits in common with spatial regions (Casati and Varzi, 1994) but can be distinguished from them in that they are immaterial parts of organisms. Just like material continuants, they move with the movement of their bearers from one spatial region to another.

Biology occupies itself mainly with three-dimensional continuants, which typically have tops and bottoms, insides and outsides, fronts and backs. Processes, by contrast, have a beginning, middle, and an end. In contrast to continuants, processes unfold themselves along a temporal axis so that, for example, your childhood and your adulthood are temporal parts of the process that is your life.

Here, the concern is with two different, complementary perspectives on the same reality: one space- and matter-oriented, the other time- and change-oriented. There are certain logical and ontological relations between these perspectives, which we make explicit in our treatment of the relations by taking account of spatial regions and points in time (see Smith and Grenon, 2004). It follows from our approach, which recognizes a radical distinction between continuants and occurrents, that there are limitations on which sorts of entities can serve as relata in any given relation. For example, it is incoherent to form a relational statement of the form *P is_a C*, because the subsumption relation cannot obtain between entities from incompatible categories. By contrast, the *has_participant* relation (as in *apoptosis has_participant cell*) requires that the first-named entity be a process universal and the second a continuant universal, rendering this relation something like a bridge between our two perspectives on reality.

Also, we need to distinguish between two types of relations at the level of instances: there are relations between continuants (for example: Mary's uterus is an instance level part of Mary), whose representation must contain a reference to points in time, and there are relations between processes, for whose representation this is not required. (The course of Mary's pregnancy is a part of Mary's life in a time-independent sense of parthood.) Since processes unfold through time, it is as if they already contain a reference or anchorage to the temporal dimension within themselves.

The placement of the relata in the ontological sextet, effectively, establishes the limitations as to which types of relations can obtain between which relata. An instantiation relation can only obtain between a particular and a universal, a participation relation only between a continuant and a process, and so forth.

5. Primitive Relations at the Level of Instances

In order to be able to define the subsumption (*is_a*), part-whole, and participation relations at the level of universals, we must first list those relations which are not further definable and, therefore, which we view as primitive. Otherwise, we will be threatened either by an infinite regress (since each definition will require ever new, more basic vocabulary for its formulation) or a circular structure (where we effectively define an expression with the help of other expressions in which the expression to be defined already appears). The selected primitive relations should be evident, self-explanatory, and neutral with respect to the various domains of science. Hence, they are relations of the sort which obtain, not only within the field of the biology, but in any domain whatsoever. Except for the instantiation relation (**instance_of**), which obtains between an instance and a universal, all of the primitive relations obtain between instances; we can then use these primitive relations to define the relations at the level of universals.

We select the following relations as primitive, drawing in large part on the results outlined in (Smith, *et al.*, 2005):

c **instance_of** *C* **at** *t*: a primitive relation between a continuant-instance and a universal which it instantiates at a given point in time. This relation corresponds to the instantiation relation in the ontological sextet which obtains between a substance particular and a substance universal.

p **instance_of** *P*: a primitive relation between a process-instance and a universal which it instantiates independently of time. This relation corresponds to the instantiation relation in the ontological sextet which obtains between an individual process and a process universal.

c **part_of** *c*₁ **at** *t*: a primitive part-whole relation between two continuant instances and a time at which the one is part of the other.

p **part_of** *p*: a primitive part-whole relation which, independently of time, obtains between two process-instances (one is a processual part, or segment, of the other).

c located_in r at t: a primitive relation between a continuant instance, a 3-dimensional spatial region which this instance occupies, and a time at which this instance occupies this region.

p has_participant c at t: a primitive relation of participation among a process, a continuant, and a point in time. This is the inverse of the participation relation in the ontological sextet, which obtains at a certain point in time between a substance particular and an individual process.

For a human reader, these relations are relatively easy to understand. But in order to use them for computer applications, the meanings of the relational expressions must be rigorously characterized by means of axioms. Work on these axioms is not yet complete, but here are a few important ones. For the **instance_of** relation, the following axioms hold. This relation applies only to an instance and a universal, in that order. No entity can be simultaneously an instance and a universal. For the **part_of** relation, we have the following (Simons, 1987): This relation is irreflexive: no entity is a part of itself. It is anti-symmetric: if x stands in the **part_of** relation to y , and y stands in the **part_of** relation to x , then x and y are identical. It is transitive: if x stands in the **part_of** relation to y , and y in this relation to z , then x stands in this relation to z . An additive principle holds, which guarantees the existence of mereological sums or wholes. A principle of differentiation holds: if x stands to y in the **part_of** relation and if x and y are not identical, then there exists a further part z of y , which has no parts in common with x . The corresponding axioms for the instance-level part-whole relation between continuants, must be modified in such a way that they contain a temporal index.

6. *Formal Definitions of Relations at the Level of Universals*

We now have at our disposal the instruments with which we can define the relational expressions that were our original goal: those between universals.

6.1. *The Definition of the Subsumption Relation*

The *is_a* relation is often identified with the relation of set inclusion, which is well-known from mathematical set theory. Under this mistaken interpretation, the **instance_of** relation corresponds to the set-theoretical

relation of class membership. In this case a definition of $A \text{ is_a } B$ would be conceivable:

$A \text{ is_a } B =_{\text{def}} \text{For all } x: \text{if } x \text{ **instance_of** } A, \text{ then } x \text{ **instance_of** } B.$

Unfortunately, this interpretation can deliver, at most, a necessary condition for the truth of $A \text{ is_a } B$. Two arguments speak against its acceptability as a sufficient condition.

(1) This interpretation permits cases in which the subsumption relation holds not between universals or types, but between merely contingent groupings of entities, as for example in: Cell nucleus in 10ml test tube **is_a** cell nucleus. This relation certainly holds without exceptions. But a cell nucleus in a 10ml test tube is not a special kind of cell nucleus, any more than a man wearing a hat, a man not swimming, or a man in Leipzig, is a special kind of man. Certainly reasoning systems of the sorts for which ontologies are useful will need to reason with such *defined classes*, and such defined classes will be indispensable for the handling of information about scientific or clinical investigations. But defined classes are to be distinguished, nonetheless, from the universals which form the subject-matter of scientific theory.

(2) The temporal aspect receives no attention in this interpretation. One reason why ‘cell nucleus in a 10ml test tube’ does not designate a universal is that the predicate in question can both apply and fail to apply to the very same entity without the latter having changed in any way. This means that the interpretation may yield false results when applied to continuant universals, as in the already mentioned case: Adult is_a child.

We can take care of problem (1) by allowing only those type-level relations to express the sorts of genuine is_a relations which should be asserted in an ontology that correspond to statements of biological science.

Problem (2) can be dealt with by acknowledging the temporal aspect in is_a relations between continuant universals. Continuants, as opposed to processes, can instantiate various different universals while maintaining their identity over the course of their existence. Thus, we must distinguish two types of is_a relation: the is_a relation between continuants includes a temporal index, and the is_a relation between processes, which is time-independent. We define:

$C \text{ is_a } C_1 =_{\text{def.}} \text{for all } c, t, \text{if } c \text{ **instance_of** } C \text{ **at** } t, \text{ then } c \text{ **instance_of** } C_1 \text{ **at** } t.$

$P \text{ is_a } P_1 =_{\text{def.}} \text{ for all } p, \text{ if } p \text{ instance_of } P, \text{ then } p \text{ instance_of } P_1.$

6.2. The Definition of the Part-Whole Relation

Two kinds of *part_of* relation can also be distinguished at the level of universals, depending upon whether the relation obtains between continuants or processes. For continuants, $C \text{ part_of } C_1$ holds if and only if every instance of C at each point in time stands in the **part_of** relation (at the level of instances) to some instance of C_1 ; for example: *Cell nucleus part_of cell*. We thus define:

$C \text{ part_of } C_1 =_{\text{def.}} \text{ For all } c, t, \text{ if } c \text{ instance_of } C \text{ at } t, \text{ then there is a } c_1, \text{ of which it holds that } c_1 \text{ instance_of } C_1 \text{ at } t \text{ and } c \text{ part_of } c_1 \text{ at } t.$

$C \text{ part_of } C_1$ says that instances of C , whenever they exist, exist as parts of instances of C_1 . And analogously for processes: $P \text{ part_of } P_1$ holds if and only if on the level of instances every instance of P stands in the **part_of** relation to at least one instance of P_1 ; for example: *Childhood part_of life*. $P \text{ part_of } P_1$ says that instances of P always exist as parts of instances of P_1 . We thus define:

$P \text{ part_of } P_1 =_{\text{def.}} \text{ For all } p, \text{ if } p \text{ instance_of } P, \text{ then there is a } p_1, \text{ such that } p_1 \text{ instance_of } P_1 \text{ and } p \text{ part_of } p_1 \text{ hold.}$

The definitions used here have a common logical structure. Each consists of a universally quantified conditional formula containing an existentially quantified formula as part ('for every x ... there is some y ...'). This logical form we call the *all-some-structure*. It captures certain logical relations in which statements about part-whole relations stand to each other. It implies that the A and the B for example in $A \text{ part_of } B$ are treated differently. So, it cannot be concluded from *human uterus part_of human* that *human* has *human uterus* as part. For while all instances of *human uterus* are at every time at which they exist instance-level parts of some instance of *human*, it is not the case that all instances of *human* have, at every time they exist, instances of *human uterus* as parts.

6.3. *The Definition of the Participation Relation*

The primitive relation **has_participant**, at the level of instances, connects a continuant, a point in time, and a process in which the continuant is in some way involved. This relation obtains for example if a particular cell is at a particular time involved in a particular process of cell transportation. For the definition of the **has_participant** relation at the level of universals we proceed in a way analogous to the above. We thus define:

P has_participant C =_{def.} For all *p*: if *p* **instance_of** *P*, then there is a *c* and a *t*, such that: *c*₁ **instance_of** *C*₁ and *p* **has_participant** *c* **at** *t*.

Here, it should be noted that *P has_participant C* asserts merely that instances of *P* require instances of *C* as bearers. Because of the all-some structure of the definition of *has_participant*, however, it does not follow that instances of *C* are always involved in processes of a certain kind. It thus for example does not follow from *human reproductive behavior has_participant human* that all humans take part in human reproduction behavior.

6.4. *The Definition of the Location Relation*

The primitive relation *c* **located_in** *r* **at** *t* on the instance level holds between a continuant and its (unique) exact location at any given time. We can then derive a defined location relation between continuants, for example between a given cell nucleus and a given cell, as follows:

c **located_in** *c*₁ **at** *t* =_{def.} for some *r*, *r*₁, *c* **located_in** *r* **at** *t* and *c*₁ **located_in** *r*₁ **at** *t* and *r* **part_of** *r*₁.

In the relation *r* **part_of** *r*₁, *r* and *r*₁ can be conceived of as special cases of instances of continuants. This relation comprises both the relation of exact location between two continuants if *r* and *r*₁ are identical and the relation of inexact location between two continuants if *r* is a proper part (for details on the proper-part relation, see Chapter 8) of *r*₁. An example of the former relation is that between a gas and a cavity that it fills completely; an example of the latter is the location relation between a testicle and a scrotum. In this manner we arrive at a formal definition for the location relation at the level of universals:

C *located_in* C_1 =def. for all c, t , if C **instance_of** c **at** t then there is some c_1 such that: C_1 **instance_of** c_1 **at** t and c **located_in** c_1 **at** t .

7. The Logic of Relations

The inverse relation R^{-1} of a two-place relation R is defined as the relation that obtains between a pair of relata if and only if the original relation R obtains when the order of relata is reversed. It is easy to define inverse relations of the instance-level primitive relations which we have discussed. The definition of the inverse relation to *is_a*, obtaining between a universal and a universal, is trivial as well:

A *has_subclass* B =def. B *is_a* A.

Adding the *has_subclass* relation does not benefit us in increasing the expressive power of an ontology: each piece of information that can be expressed with the help of the *has_subclass* relation can be expressed with the *is_a* relation. When we move to the other relations at the level of universals, whose definitions have an all-some structure, then this is not the case. Consider, for example, the relations of parthood. At the instance level x **part_of** y is true if and only if its inverse, y **has_part** x , is true. At the level of universals, however, this is not the case. Thus, for example, although the relational statement *human testicle part_of human* is true, since every instance of *human testicle* is part of some instance of the universal *human*, there is no corresponding relation with an all-some structure which links every instance of *human* with at least one instance of *human testicle*. This should not be seen as a deficit in our definition of the relations, but rather as a reflection of the way in which reality is constructed.

Nonetheless, a *has_part* relation at the level of universals that is defined in the all-some way is useful and important, for such a relation can be used to express propositions such as *human has_part heart*. The type-level *has_part* relation we need is not, however, an inverse of the *part_of* relation, as is seen in the fact that we have:

cell *has_part* nucleus
but not:

nucleus *part_of* cell.

Corresponding to our two *part_of* relations above, we have two *has_part* relations as follows:

$C \text{ has_part } C_1 =_{\text{def.}}$ For all c and t : if c **instance_of** C at t , there is a c_1 , such that c_1 **instance_of** C_1 at t and c_1 **part_of** c at t .

$P \text{ has_part } P_1 =_{\text{def.}}$ For all p : if p **instance_of** P , there is a p_1 such that p_1 **instance_of** P_1 and p_1 **part_of** p .

In contrast to the *has_subclass* relation referred to above, the *has_part* relations thus defined bring about an increase in expressive power. They allow the representation of relations at the type level which could not be captured using *part_of* alone.

The characteristics of the relations at the level of universals with respect to transitivity, reflexivity, and symmetry are presented in the following table.

Figure 1: *Logical attributes of some formal relations.*

Relation	Transitive	Symmetrical	Reflexive	Antisymmetric
<i>is_a</i>	+	–	+	+
<i>part_of</i>	+	–	+	+
<i>has_participant</i>	–	–	–	–

8. Conclusion

Simple and rigorous definitions of relational expressions make it possible to render their meaning in a lucid way. For example, it is obvious which relation an expression such as ‘*is_a*’ stands for. Further, the form of the definition enables us clearly to see, not only the limitations on the potential relata of the relation, but also the logical properties of the relations between universals within an ontology. Our definitions are formulated in such a way that they enable a unified treatment of the corresponding

relational expressions in all biological ontologies. The methodology used to create these definitions can be used in all facets of biological science, as well as other scientific domains. In this way, it can make a contribution to the project of rendering the multiplicity of biological ontologies interoperable, even though many of these ontologies were developed for entirely different reasons and with entirely different goals. Our methodology allows for the application of automated inference mechanisms as well. Nevertheless, there is still much work to be done in axiomatizing primitive relations.

Chapter 11: Four Kinds of *Is_a* Relation

Ingvar Johansson

1. *General Introduction*

In many corners of the information sciences, including work on knowledge representation, description logics, and some object-oriented programming languages, the so-called *is_a* relation plays a prominent role. In this chapter, it will be argued that there are both material and formal reasons to distinguish between four kinds of *is_a* relation, which we will call (1) subsumption under a genus, (2) subsumption under a determinable, (3) specification, and (4) specialization.⁵¹

Genus-subsumption is the traditional way of creating classificatory trees of natural kinds; in particular, of creating the famous hierarchies of plants and animals in biology. However, it is also used in more practically oriented classifications of kinds such as citizens, patients, furniture, clothes, vehicles, and so forth. On the linguistic side, this kind of subsumption is usually mirrored, linguistically, by relations between nouns.

Determinable-subsumption, on the other hand, is not concerned with natural kinds but with qualities (properties) of different generality. For instance, as a determinate *scarlet* is subsumed under the determinable *red*. In such cases, we find relations between adjectives on the linguistic side.

Even though today, in the information and computer sciences, the two expressions ‘*a* is a *specification* of *b*’ and ‘*a* is a *specialization* of *b*’ are quite often used as synonyms of ‘*a is_a b*’, the terms ‘specification’ and ‘specialization’ will here be given more restricted meanings. In fact, these restricted meanings come close to the pre-computer world meanings of these terms. Whereas subsumptions are typically concerned with natural kinds and qualities, specifications and specializations are typically concerned with activities and processes. Prototypical specifications come out on the conceptual level primarily as relations between a verb-plus-adverb expression and a verb: ‘painting carefully’ is a specification of ‘painting’. Analogously, specializations come out primarily as relations between a verb plus a whole adverbial adjunct clause and a (usually

⁵¹ The *part_of* relation and the *instance_of* relation are not *is_a* relations at all, even though in natural languages one can say things such as ‘It *is a* part of the play’ and ‘He *is a* Swede’ (see Taivalsaari, 1996; Smith and Rosse, 2004; Smith, *et al.*, 2004).

transitive) verb which, when substantivized (converted into the language of nouns), gives: ‘painting a table’ is a specialization of ‘painting’. Of course, when the verbs are substantivized, nouns and adjectives can be used to represent specifications and specializations, too (see Figure 1).

In Figure 1, we have related some *is_a* expressions to some corresponding ordinary language sentences. In relation to this list, we will introduce a distinction between *the realist mode of speech* and *the conceptual mode of speech*, respectively. When a man in the street, or a scientist, asserts ‘a cat is a mammal’, he is talking about something that he believes to exist independently of his assertion. But when an information scientist writes ‘cat *is_a* mammal’, he may take himself to be talking only of concepts. The man in the street talks in the realist mode and the information scientist in the conceptual mode of speech; whereas the former may be said to look *through* concepts (and at the world), the latter may be said to look only *at* concepts (see Johansson, 2006). In everyday discourse, people switch from the realist to the conceptual mode when they are reading dictionaries and are reflectively translating between languages. The assertions ‘The German word ‘*Baum*’ means tree’ and ‘The German word ‘*alt*’ means old’ are assertions in the conceptual mode of speech. Each is in effect saying that a German and an English word have a concept in common. Assertions such as ‘*Dieser Baum ist alt*’ and ‘*This tree is old*’ belong to the realist mode of speech.

The distinction, now presented, has affinities with Rudolf Carnap’s classic distinction between the material and the formal mode of speech (in German: ‘*inhaltliche und formale Redeweisen*’) (1934). In fact, it can be looked upon as a version of Carnap’s distinction that has been freed from its original positivist-conventionalist setting and tied to a realist framework. The left column of Figure 1 can be read in both the conceptual and the realist mode of speech. The assertion ‘cat *is_a* mammal’ can be read either as ‘the concept of cat is a concept that is subsumed under the concept of mammal’ or as ‘the class of cats is subsumed under the class of mammals’. Note that, even though cats have (inherit) all the properties which mammals have in so far as they are mammals, the concept *cat* does not have all the properties that the concept *mammal* has. In *The Description Logic Handbook* (Baader, *et al.*, 2005, p. 5) it is asserted that:

The IS-A relationship defines a hierarchy over the concepts and provides the basis for the ‘inheritance of properties’: when a concept is more specific than some other concept, it inherits the properties of the more general one. For example, if a person has an age, then a woman has an age, too.

Figure 1: *Examples of is_a relations*

<i>is_a</i> expressions	Corresponding ordinary sentences about one individual case or about classes of cases	Corresponding ordinary sentences that are (at least seemingly) directly about universals
cat <i>is_a</i> mammal	a cat is a kind of mammal; cats are mammals	the cat is a mammal
mammal <i>is_a</i> animal	a mammal is a kind of animal; mammals are animals	
sailing ship <i>is_a</i> ship	a sailing ship is a kind of ship; sailing ships are ships	
ship <i>is_a</i> vehicle	a ship is a kind of vehicle; ships are vehicles	
scarlet <i>is_a</i> red	a scarlet thing is a kind of red thing; all scarlet things are red things	scarlet is a red hue
red <i>is_a</i> color	a red thing is a kind of colored thing; all red things are colored things	red is a color
running <i>is_a</i> activity	to run is to perform a kind of activity; all cases of running are cases of activity	running is an activity
painting <i>is_a</i> activity	to paint is to perform a kind of activity; all cases of painting are cases of activity	painting is an activity
careful painting <i>is_a</i> painting	painting carefully is a way of painting	careful painting is painting
house painting <i>is_a</i> painting	to paint a house is to paint (a certain kind of thing)	painting a house is painting
outside painting <i>is_a</i> painting	to paint an outside is to paint (a certain part of a thing)	painting an outside is painting
summer painting <i>is_a</i> painting	to paint in the summer is to paint (at a certain time of the year)	painting in the summer is painting

The quotation is understandable, but it conflates the realist and the conceptual mode of speech. Neither of the concepts ‘person’ or ‘woman’ has an age; but what can be referred to by means of these concepts, i.e., people and women, always have ages. As will be shown in what follows, in

order to become clear about the *is_a* relation in the conceptual mode of speech, one has to investigate some corresponding assertions that belong to the realist mode of speech. Sometimes, I will make it explicit when I switch between these modes of talking, but mostly I will trust that the context makes my mode of speech clear.

2. *Distinguishing between Sets and Classes*

The *is_a* relation can have as its relata real classes, concepts, or terms for concepts. Unfortunately, nowadays the terms ‘set’ and ‘class’ are often used as synonyms, but here they will be kept distinct (see Smith, *et al.*, 2005; Smith, 2005; Feibleman, 1974). Sets (as I will use this term) can be constructed by a simple act of will. Hence, a set can be created by means of artificial groupings (e.g., the set consisting of my neighbor’s cats together with his house), but no real class, such as the class of cats or the class of dogs, can be so delineated. A real class is a collection of entities that share a general language-independent feature (a universal or a type) in common. Such classes can be divided into two sorts (i) the extensions of universals (classes which consist of all and only the instances of some universal, for instance the class of all human beings), and (ii) defined (or partly fiat) classes, which are subclasses of extensions of universals delineated by means of some artificially created boundary (for instance: the class of all human beings in Leipzig).

Sets are identified by their members. The set of my neighbor’s three cats is the same as the set {Tibbles₁, Tibbles₂, Tibbles₃}, and this set remains the same even if he gives Tibbles₃ away to his daughter. Classes, in contrast, are identified by the universals and any fiat demarcations in terms of which they are defined. This means that, when considered with respect to time, classes (but not sets) can remain identical even while undergoing a certain turnover in their instances. Two distinct classes may have the same extension, but no distinct sets can have the same members. Hence, there is only one null set; but, in the sense of ‘class’ used here, the development of science forces recognition of several distinct zero-classes, i.e., classes that lack members. Famous examples of such classes from the history of science are phlogiston, planets that move around the Earth, and electron particles that orbit a nucleus.

To every non-zero class there is a corresponding set, but there is not a corresponding class for each non-zero set. In order to make this point more apparent, it will be helpful to compare subsumption schemas for classes

with those for sets. Let us take a look at a subsumption schema that consists of four levels of classes. It is only for the sake of expositional simplicity that each class in Figure 2 is divided into exactly two subclasses.

Figure 2: *A formal class subsumption schema*

Level 1	highest class: class A(1)							
Level 2	class A(2)				class B(2)			
Level 3	class A(3)		class B(3)		class C(3)		class D(3)	
Level 4	class A(4)	class B(4)	class C(4)	class D(4)	class E(4)	class F(4)	class G(4)	class H(4)

This schema for class subsumption must by no means be regarded as identical with the similar schema for set inclusion illustrated in Figure 3:

Figure 3: *A formal set inclusion schema*

Level 1	set A(1)							
Level 2	set A(2)				set B(2)			
Level 3	set A(3)		set B(3)		set C(3)		set D(3)	
Level 4	set A(4)	set B(4)	set C(4)	set D(4)	set E(4)	set F(4)	set G(4)	set H(4)

In relation to Figures 2 and 3 some of the things already said about the class-set distinction will become apparent. If none of the lowest classes of a subsumption schema is a zero class, then a corresponding set inclusion schema with the same number of sets as classes can always be constructed. One has only to regard the instances of each class as members of a corresponding set; however, the converse operation is not always possible. For example, let set A(4) be the set of cats that corresponds to the class of cats and let set B(4) be the set of red instances that corresponds to the class of red instances. The set A(3) is then simply the union of the sets A(4) and B(4), but there is no corresponding class A(3), because every class has to have some kind of internal coherence of the sort that is provided by a

common universal, but there is no such common universal between the class of cats and the class of red instances.

The following are some more examples that, like the cat and red example, will clarify the intuition behind such a conception. There is a *set* whose members consist of all temperature instances and of all mass instances, but there is no corresponding *class*. There is a set whose members consist of all molecules and of all cells, but there is no corresponding class.

In philosophy, much has been said about what, if anything, can constitute the kind of internal coherence (or unity) which distinguishes classes and sets (or, more generally, between natural kinds and arbitrary collections). The debate is still going on. The opposing positions are called realism, conceptualism, and nominalism. Realism entails that classes exist in reality independently of minds. Conceptualists hold that classes depend upon mental acts. Nominalists encourage us to believe that classes are simply that which is picked out by the use of general terms. The view put forward here might be called realist (there are completely mind-independent classes) with an admixture of conceptualism (some classes are partly fiat; more about this after Figure 5 below). However, for the purposes of this paper, it is enough if the reader accepts some conception of internal coherence that makes the distinction between class subsumption and set inclusion viable.

Let us summarize:

- from a semantic point of view, no class can be introduced or defined merely by means of an act of will, though many sets can;
- from an ontological point of view, sets can be identified with their members, but classes cannot; there can be only one zero set, but there can be many zero classes;
- from a temporal-ontological point of view, (i) certain kinds of sets can be tied to temporally located particulars; (ii) there are classes of activities and processes just as there are classes of objects and quality instances.

3. *Genus-Subsumption versus Determinable-Subsumption*

Classes of natural as well as artificial kinds (e.g., atoms, molecules, plants, animals, furniture, clothes, and vehicles) may stand in subsumption relations, but so may classes of qualities (e.g., colors, volumes, masses, and dispositional properties). As the class of cats is subsumed by the class of

mammals, which, in turn, is subsumed by the class of animals, so the class of scarlet instances is subsumed by the class of red instances, which, in turn, is subsumed by the class of color instances. With respect to individual things and spatiotemporal quality instances, these subsumptions imply:

- necessarily, if a certain particular is a cat then it is a mammal, and if it is a mammal it is an animal;
- necessarily, if there is an instance of being scarlet then there is an instance of being red, and if there is an instance of being red there is an instance of being colored;
- necessarily, if a certain particular is an animal, then it has to be an animal of a certain kind;
- necessarily, if there is an instance of being colored, then there is also an instance of some specific color hue.

Early in the twentieth century, the Cambridge philosopher W. E. Johnson (1921) argued against the view the two triple-subsumptions cat-mammal-animal and scarlet-red-color represent one single subsumption relation that relates different kinds of entities (natural kinds and qualities, respectively). Rather, he held, there are two different kinds of subsumption relation, forming genera-species hierarchies on the one hand and what Johnson termed determinable-determinate hierarchies on the other (see also Johansson, 2000). The reason for the assumption of one single subsumption relation is that both cases have, on the level of the corresponding extensions, the same class inclusion relation in common. The reason for the difference comes from the fact that species and genera (and all natural and artificial kinds of things) have monadic qualities by means of which they can be characterized, whereas determinate and determinable qualities cannot be so characterized. They are themselves qualities and so can only be characterized by means of their *similarity relations* to other qualities. The class of mammals can be defined as belonging to the *genus* animal, and as such having the specifically differentiating feature (*differentia specifica*) that the females are normally able to produce milk by means of which their offspring are first fed. The class of red instances cannot similarly be defined as colors (which would be the genus) that have in common a certain *differentia specifica* that is distinct from just being red. John Searle (1959, 143) describes this difference between species and determinates as follows:

A species is a conjunction of two logically independent properties—the genus and the differentia. But a determinate is not a conjunction of its determinable and some other property independent of the determinable. A determinate is, so to speak, an area marked off within a determinable without outside help.

If we define mammals (for short) as feeding-offspring-with-milk animals, then ‘feeding-offspring-with-milk’ and ‘animal’ are treated as being logically independent, i.e., they can neither be defined in terms of nor subsumed by each other. Even though there are no plants that produce milk, such plants are not logically impossible. One can then adequately say, with Searle, that mammals are marked off from other animals with outside help. But one cannot similarly mark off red from color (and scarlet from red) with outside help.

The need to distinguish between the genus-species distinction on the one hand and the determinable-determinate distinction on the other becomes even more obvious if we consider several subsumption levels simultaneously. Table 2 is a jointly exhaustive and pairwise disjoint (JEPD) subsumption schema that consists of four levels of classes which can be assumed to represent natural kinds of some sort. The classes on each level are mutually exclusive, and this entails that no class is subsumed by more than one class on the level above it. The schema ranges from a highest class (genus or determinable) via two intermediate levels to the lowest classes (species or determinates). All classes on the intermediate levels are species or determinates *in relation to* the higher and subsuming classes, but genera or determinables *in relation to* the lower and subsumed classes. Only the highest genus/determinable is a genus/determinable in a non-relative sense, and only the lowest species/determinates are species/determinates in this same non-relative sense.

For simplicity’s sake, we will abstract away from epistemology and ontological error, and talk as if all examples represent subsumption relations between non-empty classes. Genus-subsumption schemas represent the way pre-Darwinian biologists classified plants and animals, but such schemas are often used today outside of phylogenetic taxonomy. When a genus-subsumption taxonomy has become established, it can be used to lay down so-called Aristotelian real definitions, i.e., definitions that are primarily definitions of universals rather than of concepts. Philosophers who claim that *only* concepts can be defined are doing one of two things: they either (explicitly or implicitly) deny the existence of language-independent universals, or they restrict the term ‘definition’ in such a way that many definitions in the natural sciences cannot be called definitions.

Figure 4: *The formal structure of Aristotelian definitions of genera and species*

Class	Definition
class A(2)	$A(1) \cap a$
class B(2)	$A(1) \cap b$
class A(3)	$A(2) \cap c = A(1) \cap a \cap c$
class B(3)	$A(2) \cap d = A(1) \cap a \cap d$
class C(3)	$A(3) \cap e = A(1) \cap b \cap e$
class D(3)	$A(3) \cap f = A(1) \cap b \cap f$
class A(4)	$A(3) \cap g = A(1) \cap a \cap c \cap g$
class B(4)	$A(3) \cap h = A(1) \cap a \cap c \cap h$
class C(4)	$B(3) \cap i = A(1) \cap a \cap d \cap i$
class D(4)	$B(3) \cap j = A(1) \cap a \cap d \cap j$
class E(4)	$C(3) \cap k = A(1) \cap b \cap e \cap k$
class F(4)	$C(3) \cap l = A(1) \cap b \cap e \cap l$
class G(4)	$D(3) \cap m = A(1) \cap b \cap f \cap m$
class H(4)	$D(3) \cap n = A(1) \cap b \cap f \cap n$

In a complete system of Aristotelian definitions for any given domain, one would start from the highest genus and present, stepwise, the definitions of the lower classes until the lowest classes (species) have been defined. In each such step the subsuming class is divided into two or more subsumed classes by means of some differentiating quality requirements. The classic Aristotelian example is ‘man =_{def} rational animal’ – signifying that the subsumed class *man* is defined by means of a more general subsuming class (*animal*) plus a quality requirement, namely that the class man should have the quality *rationality* as its specific difference in relation to the other classes on the same level. The definitional route just described is also used in much computer programming. (For a formally more detailed exposition of Aristotelian definitions and of relations between species and genera, see Smith, 2005. In Figure 4, some features of importance for this paper are highlighted.)

For the sake of simple exposition, we will introduce symbols for what might be called class intersection (\cap) and class union (\cup), respectively. Though the notions of intersection and union belong to set-theory and are purely extensional, this usage is not intended to indicate that we define a class in terms of its extension. It merely means that instead of ‘man =_{def} rational animal’, we can write ‘man =_{def} rational \cap animal’; and instead of ‘red =_{def} dark red or light red’ we can write ‘red =_{def} dark red \cup light red’. Let us now assume that we have one highest genus, and fourteen quality

classes (a, b, c, ..., n), one for each *differentia specifica*. All the classes, except the highest one, can then be defined, as in Figure 4.

From a purely definitional point of view, all the classes from A(2) to H(4) become classes of natural kinds, not classes of qualities, only because the highest class A(1) is a natural kind. By definition, if the presumed specific differences do not give rise to mutually exclusive classes, they cannot be called *differentia specifica*. In definitions like these the highest genus, as well as all the species-differentiating qualities, has to be – in relation to the subsumption schema – undefined (Berg, 1983). As is easily seen in Figure 4, the lower classes have (inherit) all the qualities that are essential to the classes above them.

Aristotelian definitions are advanced in the realist mode of speech. If we switch to the conceptual mode of speech, the given definitions of the classes turn into definitions of the corresponding concepts. If, in the course of scientific development, a specific taxonomy is revised, then new real definitions have to be substituted for the old ones. When this happens, it is often the case that new or partly new concepts have to enter the scene. When, for instance, it was discovered that the class of whales should not be subsumed under the class of fishes but under the class of mammals, the concepts of both whale and fish had to be redefined (Johansson, 1986).

When Figure 2 is used to represent subsumptions under determinables such as length, color, and mass, the following should be noticed. If one wants to use the schema as a basis for definitions, one cannot proceed as in the case of genus-subsumptions. This is because: (i) trivially, one cannot create divisions of a class only by means of the class itself, and (ii) since the highest class is now a determinable, there are no qualities external to the class that can create subsumed classes. Therefore, the only way possible is to define the higher classes by means of the lower ones; which means that the lowest ones have to be regarded as undefined in relation to the schema. Since the lowest classes do not overlap, the definitions of the higher classes have to be made by means of the operation of union (\cup). The schema in Figure 2 can then be used to make the definitions stated in Figure 5.

Figure 5: *The formal structure of definitions of determinables by means of determinates*

Class	Definition
class A(3)	$A(4) \cup B(4)$
class B(3)	$C(4) \cup D(4)$
class C(3)	$E(4) \cup F(4)$
class D(3)	$G(4) \cup H(4)$
class A(2)	$A(3) \cup B(3) = A(4) \cup B(4) \cup C(4) \cup D(4)$
class B(2)	$C(3) \cup D(3) = E(4) \cup F(4) \cup G(4) \cup H(4)$
class A(1)	$A(2) \cup B(2) = A(4) \cup B(4) \cup C(4) \cup D(4) \cup E(4) \cup F(4) \cup G(4) \cup H(4)$

Some observations on the set-class distinction may be of relevance here. If the definitions given would be definitions of sets instead of classes, then it would be tautologically and vacuously true that the union of A(4) to H(4) exhausts the set A(1); but if we are dealing with classes, however, then the highest determinable has to ensure that there is an internal coherence among the lowest determinates. Otherwise the latter would not be able to be subsumed under the *class* A(1). Therefore, the definition of the *class* A(1) as the union of the classes A(4) to H(4) is in effect a statement (non-vacuously true or false) that says that the members of the classes A(4) to H(4) jointly exhaust the class A(1). In those cases where the highest determinable and the lowest determinates, but no classes in between, are naturally pre-given classes (which we think is a very important sort of case (Johansson, 2000)), then all the in-between classes can be created by means of conventions. We then get a number of partly fiat classes, for which the conventionality in question is bounded by one bona fide class at the top of the subsumption schema and many bona fide classes at the bottom.

When fiat classes of the kind mentioned are created on some given level, then one can, in principle, let these classes be either overlapping or mutually exclusive, but systems with mutually exclusive classes function much better from a communicative point of view; they simply contain more information. Then, for instance, one knows for sure that if one person says ‘this is an A(3)’ and another person says ‘this is a B(3)’, both cannot be right.

In everyday life, we divide length instances into classes such as *very short*, *short*, *medium*, *long*, and *very long*; temperature instances are similarly divided into classes such as *very cold*, *cold*, *neither cold nor warm*, *warm*, and *hot*. Classes like these can both subsume more determinate classes as well as be subsumed under even broader classes. In

physics, the same determinables ground linear scales. Such scales are special cases of determinable-subsumption. In themselves, they contain only two levels, the level of the highest determinable (length and temperature, respectively) and the level of the lowest determinates. The latter level contains (is the union of) infinitely many classes, one corresponding to each real number. For instance, the concept of 5.000789000 m refers to one class of length instances, and the concept of 74.67823000 m refers to another class. In all probability, many such classes are zero classes.

One difference between genus-subsumption and determinable-subsumption can now be summarized as follows: definitions based on determinable-subsumptions have to move bottom up with the help of the operation of class union, whereas definitions based on genus-subsumptions can also move top down with the help of the operation of class intersection.

Both kinds of *is_a* subsumption relation distinguished so far must be kept distinct from another relation that is also sometimes called ‘subsumption’, namely the relation between an individual (particular) and a class. To keep them distinct, the latter relation should be called ‘instantiation’ or ‘instance_of’. Hopefully, an example is enough to make the distinction clear. If Pluto is a brown dog, then both the statements ‘Pluto instance_of dog’ and ‘Pluto instance_of brown’ are true, but the statements ‘Pluto *is_a* dog’ and ‘Pluto *is_a* brown’ are misnomers.

4. *Specification*

Is_a relations such as ‘careful painting *is_a* painting’, ‘careless painting *is_a* painting’, ‘fast painting *is_a* painting’, and ‘slow painting *is_a* painting’ seem to conform neither to what is typical of genus-subsumption nor to what is typical of determinable-subsumption. We will call them *specifications*. Let us explain. The class *careful painting* is not identical with the intersection of two logically independent classes: *painting* and *careful*. There is no class carefulness that exists as an independent entity. Carefulness is always *careful activity of some sort*. Furthermore, the carefulness in ‘careful painting’ is distinct from the carefulness in ‘careful reading’, ‘careful cleaning’, ‘careful watching’; each of these carefulnesses is logically secondary to, and takes part of its essence from, the kind of activity that is in each case mentioned. Therefore, careful painting cannot be subsumed under painting as a species is subsumed under a genus. And what goes for *careful painting* goes for *careful painter*, too. It is a well

known fact in philosophy, linguistics, and the information sciences that (to talk in the conceptual mode of speech) the extensions for expressions such as ‘being a careful painter’, ‘being a fast painter’, and ‘being a good painter’ cannot be analyzed as the intersections of the extension of the expression ‘being a painter’ with the extensions of the expressions: ‘being careful’, ‘being fast’, and ‘being good’, respectively.

The difference between specification and determinable-subsumption is not equally clear, but one aspect of this distinction is the following. In the way we have shown, determinable-subsumption allows for definitions by means of the union of the subsumed classes, but it seems impossible to define any activity as the union of a number of specifications. For instance the extension of ‘painting’ cannot be regarded as identical with the union of the extensions of ‘careful painting’, ‘careless painting’, ‘fast painting’, ‘slow painting’, and so on for all possible specifications. Unlike genus-subsumptions and determinable-subsumptions, specifications cannot ground definitions at all.

The general remarks made above in relation to activities can be repeated in relation to processes (e.g., burning, digesting, and circulating). However, it has to be noted that some possible specifications of activities (e.g., careful and careless) cannot be specifications of processes, whereas others (e.g., fast and slow) are possible as specifications of both activities and processes. Specifications differ in structure from both genus-subsumptions and determinable-subsumptions, but it is easy to conflate them, and it is especially easy to conflate specification with determinable-subsumption. The distinction is nonetheless reflected in everyday language. We say that ‘painting is a *kind* of activity’ but that ‘painting carefully is a *way* of painting’. The crux of the matter is that different activities are not specifications but determinates of ‘activity’. That is, ‘painting’ is a determinate that is determinable-subsumed by ‘activity’, whereas ‘careful painting’ is a specification of ‘painting’; similarly, ‘careful activity’ is a specification of ‘activity’. This complication can create a need to combine in one and the same classificatory tree both determinable-subsumptions (painting → activity) and specifications (careful painting → painting).

The relation of specification seems not to be confined to activities and processes. Whereas (consciously perceived) color hues obviously are determinable-subsumed under the class of (consciously perceived) colors, the same is not true for color intensities and degrees of color-saturation. They seem to be specifications of color hues just as carefulness is a specification of activities. When two different color hues, say a determinate

red and a determinate blue, have the same intensity (or degree of saturation), the intensity (saturation) is logically secondary to, and takes part of its essence from, the color hue in question; not the other way round. The fact that color hues are determinates but color intensities and saturations are specifications is quite compatible with the fact that color hue, color intensity, and color saturation can, as in the Munsell color solid, be ordered along three different dimensions in an ordinary picture or in a three-dimensional abstract space (compare Figure 8, which combines a subsumption relation with one specification relation).

5. *Specialization*

Here are some examples of *is_a* relations that are *specializations*: ‘house painting *is_a* painting’, ‘outside painting *is_a* painting’, ‘summer painting *is_a* painting’, ‘car driving *is_a* driving’, ‘food digesting *is_a* digesting’, and ‘paper printing *is_a* printing’. In these cases, the class on the left of the *is_a* relation does not *specify* the activity mentioned on the right; it is doing something else. It relates the activity mentioned on the right to something (houses, outsides, and summers) that exists completely independently of this activity. This fact makes it immediately clear that specializations cannot possibly ground definitions of the activities that they are specializing. As we normally use the concept of specialization, we can say that one painter has specialized in painting houses and another in painting chairs, one in painting outsides of houses and another in painting insides. This is our main reason for having chosen the label ‘specialization’. However, our choice is in conformity with the terminology of a paper that has previously mentioned the feature that I am now trying to make even more clear; the author in question talks about *specializing criteria* as a certain kind of subsumption (*is_a*) principle (Bernauer, 1994).

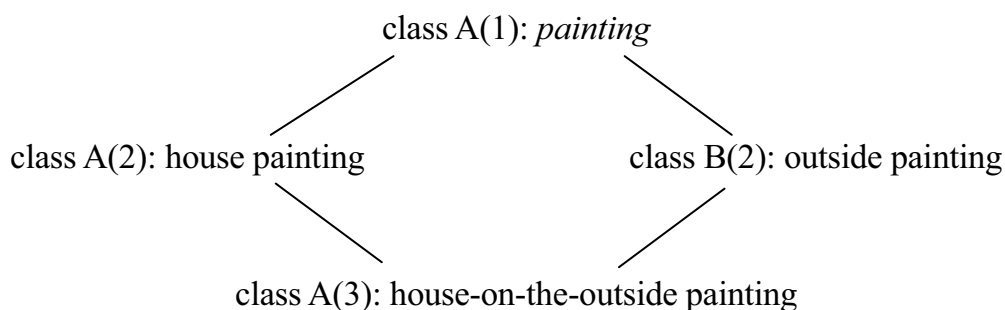
Some activities are simply activities performed by a subject (e.g., swimming and running), whereas others also involve one or several objects that are acted upon (e.g., painting a house and driving a car). Similarly, some processes simply occur in an object (e.g., rusting and burning), whereas others involve also one or several objects that the process in question acts upon (e.g., digesting food and printing papers). It is only in the *acting-on* kind of cases that specialization of activities and processes can come about. When there is talk about painting, driving, digesting, and printing as such, one knows that there is an object that has been abstracted away. It is this missing object that re-enters when a specialization is

described, or when a corresponding *is_a* relation is stated. Nothing like this occurs in subsumptions and specifications.

In all of the examples used above, the specializations are described by means of transitive verbs (or substantivizations of such verbs). And this is no accident. Transitive verbs are defined as verbs that can take (and often require) an object, whereas intransitive verbs cannot. Nonetheless, even intransitive verbs admit of specializations. This happens when the activity (process) in question becomes related to a certain kind of time period or a certain kind of place: ‘summer swimming *is_a* (specialization of) swimming’ and ‘pool swimming *is_a* (specialization of) swimming’.

Normally, an activity can be specialized in several different directions. One can paint a house, a car, or whatever. As soon as the object painted is such as to have both an outside and an inside, one can paint either the one or the other. Similarly, one may paint at a certain time of the year or at a certain kind of place. Therefore, some specializations have to have more than one *is_a* relation to the next level. Let us specialize ‘painting’ along two different directions: what kind of object that is painted and which part of an object that is painted. This, then, yields the following *is_a* schema:

Figure 6: *A double-specialization schema*



In words: house-on-the-outside painting *is_a* house painting; house-on-the-outside painting *is_a* outside painting; house painting *is_a* painting; and outside painting *is_a* painting.

6. *Single and Multiple Inheritance*

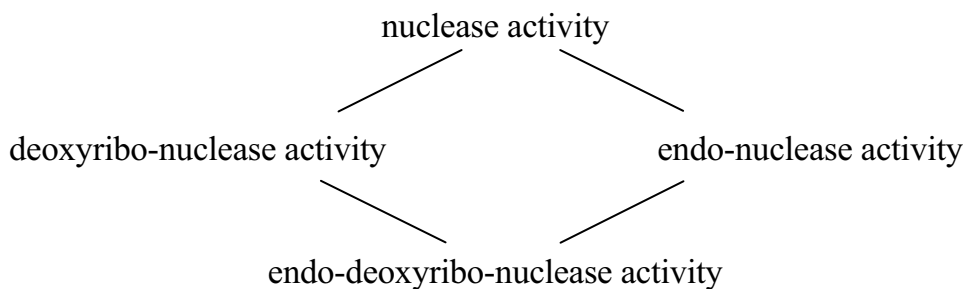
In the distinction between single and multiple inheritance, the concept of inheritance seems originally to have referred to inheritance of qualities in genus-subsumptions. A subsumed genus inherits all the properties that are essential to the subsuming classes. If a certain genus is subsumed by only one class on the nearest upper level, then there is single inheritance of qualities. If it is subsumed by two or more genera, then there is multiple inheritance. In determinable-subsumptions there are no real inheritances of

qualities apart from the inheriting of the highest determinable; the rest is, as we have explained, a matter of mere unions of the lowest determinates. Specifications, too, do not involve any literal quality inheritances. Nonetheless, the distinction between single and multiple inheritance is sometimes applied to each kind of *is_a* relation that we have distinguished. This means that when, in this general sense, it is stated that there is multiple inheritance, it is merely stated that the left-hand class of an *is_a* relation has some *is_a* relation to more than one class on the next level up.

From our remarks on genus-subsumption and determinable-subsumption, it follows that in both cases the default norm for such *is_a* hierarchies should be that they contain no multiple inheritances. With respect to specification, however, it does not even make sense to speak about multiple inheritance *of only specifications*. As we have analyzed ‘careful painting’, it can only have a specification relation to ‘painting’, since ‘careful’ in ‘careful painting’ has no complete meaning independently of painting. With respect to specializations, however, things are completely different. Here we get multiple inheritances as soon as there are two or more different directions that a specialization can take. In Figure 5, ‘house-on-the-outside painting’ is multiply (doubly) inherited.

Multiple inheritances are consciously and, according to my analysis, correctly used in the Gene Ontology (see *Gene Ontology*). The Gene Ontology Consortium states that ‘GO terms are organized in structures called directed acyclic graphs (DAGs), which differ from hierarchies in that a “child” (more specialized term) can have many “parents” (less specialized terms)’ (Gene Ontology Consortium, ND). The GO consortium uses the concept of ‘specialization’ as a synonym for ‘*is_a* relation’, but whenever a child in a GO graph has more than one *is_a* parent, then at least one of the *is_a* relations in question is a specialization in my restricted sense.

Figure 7: A specialization schema with examples from the Gene Ontology

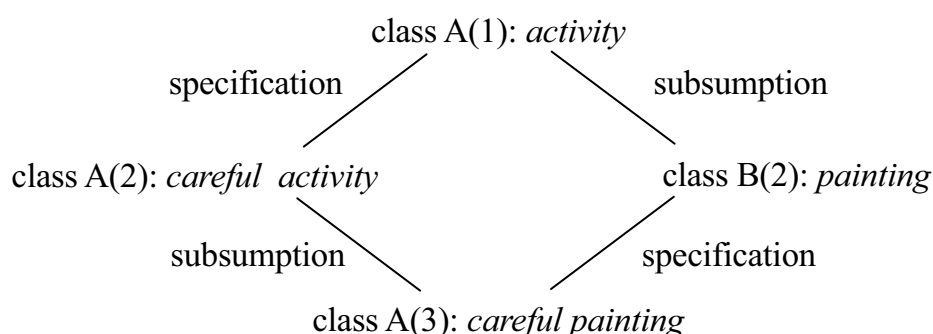


Let us exemplify. In the GO ontology for molecular functions one finds (hyphens added) ‘endodeoxyribo-nuclease activity’ (GO:0004520) inherited from both ‘deoxyribo-nuclease activity’ (GO:0004536) and ‘endo-nuclease activity’ (GO:0004519); both of the latter are, in turn, inherited from ‘nuclease activity’ (GO:0004518). Setting the arrows of GO’s graphs aside, these specializations can be represented as in Figure 7.

This is one of numerous examples of specializations that can be extracted from the Gene Ontology. A nuclease activity is an activity (performed by an enzyme) that catalyzes hydrolysis of ester linkages within nucleic acids. Such activity can be specialized along at least two different directions: (i) according to *what* is acted on (deoxyribonucleic acid, DNA, or ribonucleic acid, RNA), and (ii) according to *where* the action takes place, i.e., cleaving a molecule from positions inside the molecule acted on (‘endo-’), and cleaving from the free ends of the molecule acted on (‘exo-’), respectively. Since nothing stops the specialization from going in both directions at once, we get the schema for ‘nuclease activity’ in Figure 7, which is completely analogous to the schema for ‘painting’ in Figure 6.

Specializations allow and often require multiple inheritance. They differ in structure from genus-subsumptions, determinable-subsumptions, and specifications. However, we have so far spoken of hierarchies or graphs consisting of only one of these kinds of *is_a* relation, but the different *is_a* relations can also be combined with each other. (They can also, as in the GO, be combined with the *part_of* relation.) And in such mixed cases, too, multiple inheritance can be the normal and the required kind of inheritance. Two examples may show what we mean.

Figure 8: *A combined specification and determinable-subsumption schema*

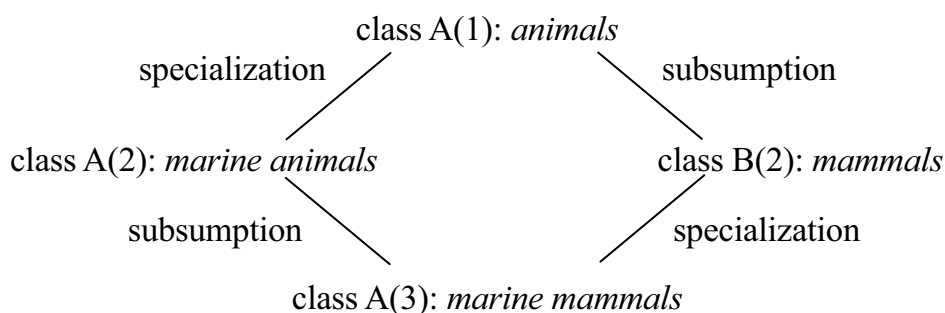


In the first example, Figure 8, ‘careful painting’ is doubly inherited. On the left hand side, the *is_a* relation is one of subsumption, but on the right

hand side it is one of specification. When the whole Figure is taken into account, symmetry is displayed. Two specifications are diagonally opposed, and so are two subsumption relations.

In traditional non-phylogenetic classifications of animals the *differentia specifica* are (broadly speaking) properties inhering in the organisms. But, of course, one can also try to classify animals according to where, when, and on what they perform various activities. Some live on land and some in the sea; some sleep during the night and some in the day; some eat meat and some do not. Therefore, classes of animals can via their activities also be made relata in specialization relations. Most mammals live on land but whales live in the sea. We may speak of a class *marine mammals* that can be placed in an *is_a* schema such as that of Figure 9.

Figure 9: *A combined specialization and genus-subsumption schema*



We have sometimes mentioned the ‘instance_of’ relation (e.g., at the end of section 2). Now, in order to avoid all misunderstandings, we need to do it again. Everything that has been said about multiple inheritance above relates to *is_a* relations and not to ‘instance_of’ relations. Trivially, an individual can instantiate many classes and in this special sense have multiple inheritance (better: ‘multiple instantiation’) when placed in a slot in a matrix. Many matrices that are used in the social sciences and in epidemiology to display correlations have this character. A simple but fictive example that contains this kind of multiple (double) inheritance for a group of hundred persons is presented in Figure 10 (from Asplund, 1968).

Figure 10: *A correlation matrix relating political views to political interest*

	Republicans	Democrats	Independents
High	14(persons)	16	5
Medium	19	17	7
Low	5	5	12

Here, each of the fourteen individual persons in the upper left slot inherits two features: having high political interest and being Republicans. Such tables must by no means be conflated with those in Figures 6 to 9 above. Note, though, that if an individual is an instance of a certain class, then he is automatically also an instance of all classes that subsume this class.

7. *Philosophy and Informatics*

Can the taxonomy of *is_a* relations presented in the above be of any use in informatics? Let us answer by means of a detour.

No observations can be reported, and no reasoning can take place, without classification. But a classification is not necessarily a taxonomy: it need not be a *systematized* classification. During medieval times, alchemists made extensive classifications of substances, and herbalists did the same with respect to plants, but in neither case was a real taxonomy created. In a sense, the alchemists and the herbalists were too practically minded. But, things changed with the advent of modern chemistry and botany. Remarkable taxonomies appeared with remarkable repercussions on scientific development.

Today, information scientists help other scientific disciplines, as well as practical endeavors of all kinds, to systematize their respective classifications. But, curiously enough, they nonetheless seem to have no deep impulse to systematize their own use of various kinds of *is_a* relations and different kinds of definitions. Despite being a philosopher by trade, I dare to assert that at least some information scientific work can be done more efficiently if those involved would (a) accept that there is a distinction to be drawn between sets and classes, and (b) become aware of the tetrachotomy of *is_a* relations that we have put forward. Thus far, much that has been created, on a purely pragmatic basis, has rested on principles that were made explicit only later. However, once discovered, such principles can be consciously put to use and, thereby, make future similar work simpler and more effective. Without any explicit talk of that special kind of *is_a* relation which we have called specialization, the authors of the Gene Ontology chose to work with directed acyclic graphs instead of the set-theoretical inclusion relation, but this fact is no reason not to make the next generation of information scientists aware of the existence of the different kinds of *is_a* relations thereby involved.

Chapter 12: Occurrents

Boris Hennig

In this chapter, we distinguish occurrents from entities of other sorts. Then, within the class of occurrents, we introduce several other distinctions that will yield a taxonomy of temporal entities. We distinguish temporally extended from instantaneous occurrents, and a further distinction is drawn between processes which have an internal temporal structure, and other temporally extended occurrents which are internally unstructured. It turns out, however, that such distinctions only apply directly to *types* of occurrents, and only indirectly to their particular *tokens* or instances. Therefore, we have to consider the way in which types of occurrents are related to their instances. It will become apparent that individual occurrents may instantiate more than one type simultaneously, where the types involved are systematically related to one another.

1. *Some Things that are not Temporal*

At first glance, it may seem that everything that exists is also temporal in some sense. Therefore, it may not make much sense to distinguish temporal from non-temporal entities. However, there are at least three types of things that may be said not to be temporal in a strict sense.

First, there are things that are prior to all temporality in the sense that they are more fundamental than everything that is temporal. For instance, if something is temporal by virtue of being or occurring in time, then time itself is either not temporal or it occurs immediately in itself. Since it is hard to make sense of the latter, time itself does not appear to be temporal. It is, rather, prior to all temporality. Accordingly, time is not an occurrent. A second class of things that are not temporal in a stricter sense consists of abstract entities, including numbers, geometrical shapes, and universals, such as the types under which concrete temporal and non-temporal things fall. These entities are also not the immediate topic of the present chapter. Types of temporal entities, however, will play a role later in this chapter, and therefore it will be good to briefly clarify the status of such types.

A type is something with respect to which concrete things may be called either typical or atypical.⁵² Types are specified by characterizing their

⁵² Compare the use of the term ‘type’ by C. S. Peirce, who introduced the *type / token* distinction into philosophy. Peirce also calls types ‘legisigns’ and thereby indicates

typical instances. For example, in order to specify the type ‘beaver’, one will have to mention the characteristic features of the instances of this type. For instance, beavers have 20 teeth. This does not mean that all or even most beavers have these features, since it may well be the case that all existing instances of a type happen to be deformed, unhealthy, or atypical in some other sense. They might all have lost one tooth. Further, a type may be realized by several different subgroups of still typical instances: for example, there are male and female exemplars of the type ‘beaver’, and there are at least three rather different ways of being a typical ant.

In the present context, we will only deal with types of *temporal* entities, such as the type ‘gastroscopy’. A concrete gastroscopy – the examining of the inside of patient Chen’s stomach with an endoscope – is what it is in virtue of instantiating this type. As its instance, a concrete gastroscopy is subject to certain rules that determine its typical and proper form. These rules do not only apply to one particular gastroscopy, but they describe its more general type, which may be instantiated by any number of instances. They state how gastroscopies are to be performed *in general*. When a doctor explains to a patient what will happen during the course of an impending gastroscopy, he is specifying these rules. By looking at the rules, one may determine what belongs to a typical gastroscopy and what does not; but this does not mean that every proper instance of the type ‘gastroscopy’ must be a typical one. Particular gastroscopies may violate the laws of typicality that apply to their type, without ceasing to be what they are. Although the type ‘gastroscopy’ is a type *of* something that occurs in time, it is not *itself* something that occurs at any time. What occurs is in each case one of its concrete instances. Thus, in this sense the type ‘gastroscopy’ is not temporal: it does not occur.

A third class of entities that are not strictly temporal consists of continuants, that is, concrete things and their properties. Things and their properties may change over the course of time, but they do not occur in time. The class of things and their properties will be taken to include concrete things, such as a particular endoscope, but also physicians, digestive systems and their parts, and such entities as the form of an endoscope, the license of a physician, the price of a medication, and the condition of a patient.

All of these things may change, which might incline one to say that they exist in time. But when continuants change, they also appear to persist

that types are specified by stating *laws* according to which their instances are to be classified as typical or not (Peirce, 1998).

through time in a way in which occurrents do not. It is of course perfectly acceptable to say that an occurrent, such as a gastroscopy, *lasts* for ten minutes, but this is not the same as when a continuant persists for ten minutes. It is difficult to draw this distinction in a clear and meaningful way. One is inclined to say that when continuants change, they also must remain the same in a certain respect. For instance, if the patient did not at least remain a human being, we would not say that *her* condition has improved.⁵³ On the other hand, one may also say of a particular gastroscopy that in some sense, it remains the same throughout its occurrence. Hence, remaining the same does not mark off continuants from occurrents. Since any more rigorous and detailed account of the distinction between continuants and occurrents would occupy too much space, we must here rely on an intuitive distinction. Continuants are things that may be said to come into being, perish, and persist throughout a period of time; as opposed to occurrents which may be said to start, end, and last for a certain time.

In this chapter, our concern is with entities that do not belong to any of the three categories just outlined. We will not discuss time itself, nor will we be concerned with abstract entities, types, or universals as such, nor will we consider such things as an endoscope, a patient, her condition, or a license. The entities that will be discussed here are entities such as the *improvement* of someone's condition, the *performance* of a gastroscopy, the *loss* of a license. These are entities that one may call temporal in a stricter sense: they happen or occur in, or over the course of, time. We will call them *occurrents*. The first question to ask is: How many general kinds of occurrents are there? In developing an answer to this question, we will gain a clearer insight into the features that distinguish occurrents from other kinds of entities. Occurrents may happen or occur at a certain time. An endoscope does not happen or occur, but its *use* or *modification* does.

2. Things that may Occur

2.1. Instantaneous vs. Extended Occurrents

A first distinction that may be drawn within the class of occurrents is that between *instantaneous* and *temporally extended* occurrents. Instantaneous

⁵³ In general, if something were to change in all possible respects, there would be nothing which would be the subject of change. See Aristotle, *Physics* I 7; Kant, 1781, B 225ff.

occurrences happen, or occur; but they do not stretch over a time interval. For instance, the very end of a gastroscopy, in distinction to its concluding *phase*, is not temporally extended. As the end of an occurrence, it occurs instantaneously. It may well be that every instantaneous occurrence is necessarily a part of a temporally extended occurrence, such as the beginning, middle, end, beginning of its last third, etc. Non-extended, instantaneous occurrences will then be nothing but the boundaries of extended occurrences.⁵⁴ This would indicate that temporally extended occurrences, which we mainly discuss in the following, are ontologically more fundamental.

2.2. Occurrences with a Generic Structure

Temporally extended occurrences are occurrences that take their course over a stretch of time. Such occurrences unfold in time; they may be said to consist in a sequence of stages, including at least a beginning and an end. Hence, one may be inclined to distinguish extended occurrences that typically take a specific course from others whose course is entirely undetermined. But are there occurrences whose structure is entirely undetermined? If this were so, we would in any case have no general names under which they would fall, since every name would associate the occurrences to which it applies with a specific type. We have already argued that every instance of a type is subject to certain standards of typicality. Hence, for all temporally extended occurrences that instantiate a type, there will be certain rules of typicality that govern their general structure.

One may object, first, that it is perfectly possible that something that happens here and now does not have any discernible structure. Hence, it would seem that there are occurrences that do not follow any specific course and to which no standards of typicality apply. But once we begin to spell out what it is that happens in this supposedly entirely unstructured way, we thereby, also, begin to determine its structure. In fact, it is already enough to call the item in question something which is occurring presently, since this already has implications regarding its structure. Any present occurrence must not yet be over; and in order to know whether an occurrence is over or not, we need to know under what conditions it would be over. In order to know this, however, we need to know the general type and structure of the occurrence in question.

⁵⁴ See Aristotle, *Physics* VI 3.

Second, it could be argued that there is a most general type of occurrent, namely the type ‘occurrent’, which we may apply to something without implying any details about its more specific structure. Hence, it seems that we may indeed refer to entirely unstructured temporally extended occurrents. This however is not true, since everything that instantiates this most general type ‘occurrent’ also necessarily instantiates a more specific type. Again, since extended occurrents are items that have a beginning and an end in time, there must be some more specific subtype for everything falling under the most general type ‘occurrent’, which at least determines the specific conditions under which this kind of occurrent may be said to have begun or have ended. When we refer to something as an occurrent without characterizing its structure, we have nonetheless implicitly claimed that it has some such structure.

Third, it could be argued, for instance, that the persistence of a certain unchanging condition need not take any structured course. As long as the condition persists, no change occurs, and nothing specific happens. Now what we have here will nonetheless be an occurrent. Hence, it will be an occurrent that takes no specific course. But even in such cases, there must be at least two things we know about it: we must be able to tell under what conditions such an occurrent begins and under what conditions it ends, and this is already enough of a typical structure.

There are some occurrents, called *energeiai* by Aristotle, which are special in that they may already be complete while they are still occurring.⁵⁵ For instance, when someone knows or sees something, he may also have known and seen it before. It would be wrong to say that he ‘is knowing it’, as if one could be engaged in an activity called knowing for some time and later be done with it. Put differently, in contrast to a gastroscopy, knowing something is not directed at a point when it will be both complete and over. Rather, to know something is already to have reached the relevant state of completion. We should distinguish between the *completion* of an occurrent and its *end* or *being over*. In the case of *energeiai*, the conditions under which they are complete differ from the conditions under which they are over. All other occurrents are also over when they are complete. Nonetheless, to know something is not a state (which is a continuant) but something that we do: it is an occurrent. Hence, there seem to be occurrents that take no specific course, since they are already complete when they occur.

⁵⁵ Aristotle, *Metaphysics* IX, 1048a18–b34.

But in any case, *energeiai* are also not entirely unstructured. Insofar as knowing, for instance, occurs during a certain time, this occurrence will also have a general structure. It must at least be determined under what conditions the possession of this knowledge may be said to take place, so that we may judge what it takes for someone to come to know and to cease to know. *Energeiai* such as knowing differ from other occurrents only in that (1) during their occurrence, the only further development they may take is to end, and that (2) this end can hardly be called their completion. (To cease to know is not to complete one's knowing.) All this means that we cannot assert that there must be occurrents without any general structure. The opposite is the case: every occurrent is an instance of some specific type.

2.3. *Internally Structured Occurrents*

The specific structure of an occurrent need not be known in detail; on the contrary, much may remain unspecified. The structure of an occurrent is already determined as soon as there are some criteria according to which it may be identified as an instance of its type. For that to be the case, we do not need to know many details.

There are, roughly, two degrees to which an occurrent may be structured. First, it may only be determined under which circumstances occurrents of its type begin and end. For instance, when someone looks for a pen, it is determined when this occurrent is complete and over, but it is not determined how long it will take and what steps, in which order, will be required. Looking for a pen has no *internal temporal structure*. If an occurrent does not have an internal temporal structure, there is no way to determine to what extent it is complete as long as it is still occurring. Before looking for a pen is over, it is not in general possible to tell how long it will take or how much of it has been done. This is also true of the persistence of an unchanging condition.

Second, there are occurrents with an internal structure. A gastroscopy, for instance, has an internal structure, since one may roughly say, at every one of its stages, how much of it is already over and what remains to be done.

One might object that in some sense, of course, we may know about every concrete occurrent to what extent it is over. In order to determine that, it seems, we only need to measure how long the complete occurrent takes and then calculate how much of this time is left. And since every

concrete occurrent must have a determinate duration, it will thus turn out that every concrete occurrent has an internal temporal structure.

But the distinction between occurrents with and without an internal temporal structure should be drawn more carefully. We want to say that an occurrent has an internal structure only if it is possible to determine how much of it has already occurred *while it is still occurring*. This holds true in the case of gastroscopies: we may determine, say, when half of it is done, even if we do not know how long the complete operation will take. That we are able to determine *post hoc* how much of an occurrent had occurred, after knowing how long the complete instance actually took, does not mean that it had an internal temporal structure in this sense. For this reason, we should rather speak of a *typical* internal temporal structure. We know how long gastroscopies typically take and what steps they typically involve, and only this enables us to tell how much of one of them is complete when it is still going on. By contrast, we cannot tell how long looking for a pen usually takes; we can therefore only say how long half of one of its instances would have been when it is over. That is, looking for a pen does not have a typical internal temporal structure.

Whether an occurrent is half over in this sense has nothing to do with its concrete, exact temporal duration. It may well be that the second half of a soccer match takes longer than the first, but we still call it a half. The typical internal structure of an occurrent is not measured in seconds, but consists in a more or less flexible sequence of steps. The distinction between occurrents with and without internal temporal structure is properly applicable only to *types* of occurrents. No particular occurrent has ever occurred that would not have had some particular internal structure. Conversely, one may say of any ongoing occurrent that its concrete internal structure is still undetermined: it is not yet established how this particular instance will in fact turn out to be structured. However, there is an interesting difference between gastroscopies and other occurrents such as looking for a pen, and this is the difference we want to point out here. Gastroscopies are of a type such that their structure is roughly determined before they are over, whereas looking for a pen is not of a type such that it would be even roughly determined in advance how long each instance takes. The rules that determine how to perform a gastroscopy also specify the typical internal temporal structure of a gastroscopy.

2.4. *Telic and Atelic Occurrents*

For all occurrents with a typical structure, whether internal or not, it is determined under what conditions they are complete or over. One may therefore want to call structured occurrents *telic*, since ‘*telos*’ means *end* in both of our two distinguished senses: the state of being complete, and the state of being over. Accordingly, one may wish to call those occurrents that have no structure *atelic*.⁵⁶ However, this distinction is not of much use, since we have already shown that there are no entirely unstructured occurrents. Being an occurrent already implies structuredness.

Antony Galton (1984, 66) defines telic occurrents as occurrents that may be interrupted before they are over. However, as we shall presently see, every occurrent that occurs here and now is, by necessity, not yet over (because, when an occurrent is over, it is no longer occurring). Therefore, one may interrupt every occurrent before it is over. There are some occurrents that cannot be interrupted before they are *complete*, and perhaps this is what Galton means. In this case, the distinction he draws would coincide with the distinction between *energeiai* and other temporally extended occurrents. *Energeiai* are complete before they are over: they can be complete in that nothing belonging to them is undone, yet they can still be going on. For instance, someone can completely see, know, or enjoy something before she actually stops seeing, knowing, or enjoying it. Hence, *energeiai* may be interrupted before they are over, but perhaps not before they are complete. But this again is not a good reason for using the terms ‘telic’ and ‘atelic’, since these may as well be taken to refer to the completion of an event, and it would be misleading to say that *energeiai* are occurrents that are *atelic* in the sense of being incomplete or not allowing for a state of completion. Rather, *energeiai* are necessarily already complete whenever they occur.

Zeno Vendler distinguishes between *accomplishments* on the one hand, which may be interrupted before they are complete, and *activities* on the other hand, which may not (see Vendler, 1972). According to Vendler, running is an activity, whereas running a certain distance is an accomplishment, and reaching the end of this distance is an achievement. Activities, in this sense, have already been going on when they are going on: whenever I *am* running, I also have run immediately before. By

⁵⁶ See Comrie, 1976, section 2.2.; also Dowty, 1991, section 2.2.

contrast, it is not the case that whenever I am running 10 meters, I also have run 10 meters immediately before.

Accordingly, moving the endoscope would be an activity, but carrying out a gastroscopy would be an accomplishment. But since, at certain times, gastroscopies consist in movements of endoscopes, this is not an ontological distinction between two kinds of occurrents, but only a distinction between different ways of referring to one and the same occurrent (see Gill, 1993). What the physician is doing when she is performing a gastroscopy may be regarded as something that is not yet over, namely a gastroscopy, or it may be regarded in abstraction from its outstanding end, namely as handling an endoscope. But the physician is not doing two things at once. The distinctions drawn by Galton and Vendler concern only different ways of referring to the same occurrent, and not even to different kinds of occurrents.

2.5. *Completion vs. End*

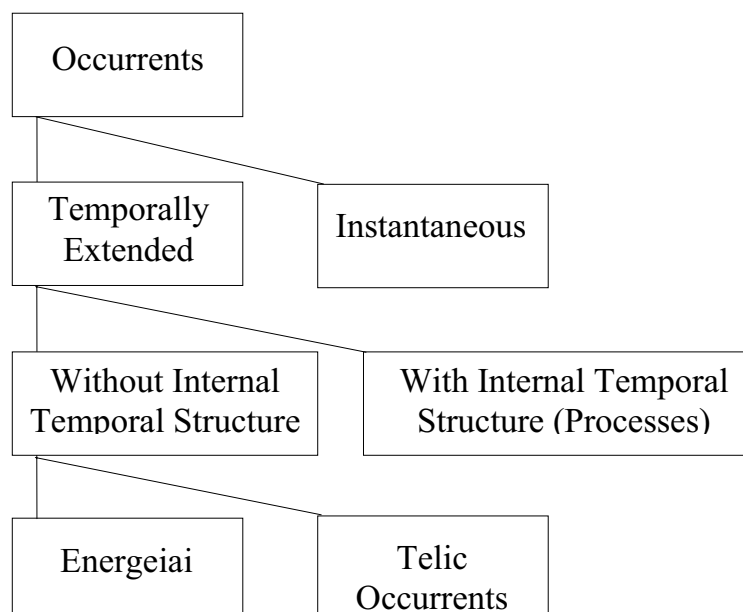
Moreover, one should not suppose that all occurrents with a generic structure are characterized, mainly, by having a certain *result* or even *goal state*. This also seems to be implied by calling them telic, but it is not true in most cases. It is true that the last step in a typical gastroscopy is the removal of the endoscope, and that the gastroscopy should normally also be completed by this step. This however does not mean that removing the endoscope is its goal: the removal of an endoscope alone is certainly not a gastroscopy, and its more important steps happen before this. Likewise, my reading a book is complete and over when I read the last page; but, of course, only when I have also read enough of the rest. In most cases, in order to specify what it takes for an occurrent to be complete, we need to mention more than its result state; we need to mention everything that typically belongs to the occurrent. Hence, the *telos* (completion) of an occurrent type always involves everything that belongs to the typical course that its instances take, not only the last step. Occurrents other than *energeiai* are over when they are complete, but not necessarily complete when they are over. As Anscombe says, ‘A man can be doing something which he nevertheless does not do’ (1957, §23). For instance, someone who is hit by a car while crossing a street *was* indeed crossing the street but *did not* in the end cross it. In the case of *energeiai*, the *telos* has nothing at all to do with the last step of the occurrent in question.

Some occurrents are indeed complete when and only when a certain result state is reached, and the *telos* of these occurrents may be justifiably identified with the state of the world that results from their occurrence. For instance, looking for a pen is complete if and only if the pen is found (or we give up the search), and it is irrelevant where and for how long the search was going on. When someone is looking for her glasses while wearing them, it is possible that no apparent searching behavior whatsoever is involved. In the present context, we call such occurrents *telic* that are complete when only their result state is reached. Such occurrents constitute an exception from the general rule, since the completion of most occurrents involves more than only reaching a result state.

2.6. *Interim Statement*

Let us briefly review the distinctions introduced so far. First, there are instantaneous occurrents and temporally extended ones. All temporally extended occurrents have a rudimentary typical structure since, for all of them, it must at least be clear under what conditions they begin and end. Some temporally extended occurrents do not possess any further internal temporal structure. These are, first, the ones that have been called *telic* at the end of last section: occurrents whose completion does not require more than the reaching of a certain end state. Second, those occurrents that Aristotle called *energeiai* also belong to this group. Since *energeiai* only occur as long as they are complete, and there is no further goal that they reach during their occurrence, they have no further internal structure. Besides these two kinds of occurrent without internal structure, there are occurrents that have a typical internal temporal structure. We will call these occurrents ‘processes’.⁵⁷ The discussion so far yields the following taxonomy:

⁵⁷ This term, ‘process’, is used in a variety of ways in the literature. Mourelatos (1978) defines processes as *atelic* occurrents; but he later withdraws this definition (1993). Stout (1997) identifies processes with types of occurrents, on which we will comment later. The Gene Ontology has used the term to designate complex, internally structured occurrents and calls simple extended occurrents ‘functions’ (Function Ontology Rules, *Gene Ontology*). Etymology would suggest that a process is something that involves a change, may be counted, and thus takes a typical course. What we call process in this paper has also been called *kinesis* by Aristotle, and by others *achievement* (Ryle, 1940, 130), *performance* (Kenny, 1963, Chapter 8), *accomplishment* (Vendler, 1972, Chapter 4), or *development* (Mourelatos, 1978).



2.7. Simple and Complex Occurrences

A further distinction remains to be discussed: that between simple and complex occurrences. Since instantaneous occurrences do not have temporal parts, one might think that they cannot be complex. Yet this is not correct. It may very well be essential to the success of a certain operation, for instance, that a physician performs two movements at once. The exact moment in which such a double movement succeeds will be an instantaneous occurrence that will nonetheless be complex, since it incorporates at least two occurrences of a different kind.

For the same reason, there may be complex temporally extended occurrences without an internal temporal structure. There may be occurrences, for instance, which typically have certain kinds of occurrences as components, but for which it may nonetheless be undetermined in what order they occur. Such occurrences have no *temporal* internal structure, since it will not be clear during their occurrence how much of them will already be over and how much remains to occur.

Are there, on the other hand, temporally extended occurrences which are simple? We have seen that all temporally extended occurrences incorporate at least two instantaneous occurrences: their beginning and their end. But this does not make them complex in the sense to which we appeal here. Here, 'complex' is applied to those temporally extended occurrences in which what happens *between* their beginning and the end, not counting the beginning and end themselves, is complex. That is, in a first approximation, we may call a temporally extended occurrence complex if it

can be broken down into further *temporally extended* occurrents, and simple if it cannot. But still, there is a sense in which every temporally extended occurrent may be divided into temporally extended parts, for every one of them can be split up into halves, thirds, etc. For this reason, it will be better to say that an extended occurrent is complex only if it can be divided into further temporally extended occurrents of *different, more specific types*. A uniform movement, for instance, can only be broken down into further uniform movements, and these are occurrents of the same type. It will accordingly count as simple. Also, waiting for an idea is a simple occurrent, since it is not known what is typically involved in this kind of occurrent other than lacking the idea, waiting for a while, and then (if you're lucky) having it. Every extended part of it is also an instance of waiting and not of any more specific type. In contrast, a gastroscopy will count as complex because it may be divided into steps that instantiate different, more specific types.⁵⁸

Thus, there are complex instantaneous occurrents, complex extended occurrents, and simple internally unstructured occurrents. But are there simple extended occurrents? A simple process should be composed of no further processes, but one should still be able to tell, as long as it is occurring, how much of it is complete. Dividing a complex process into its temporally extended parts will indeed yield such simple processes. Consider the movement that a physician makes when she inserts the endoscope into the esophagus. This movement is uniform; that is, all its extended parts are further movements of the very same type. Nonetheless, we may tell at each of its stages how much of it remains to be done, and therefore, what we have here is a simple, internally structured occurrent.

It follows that the distinction between simple and complex occurrents is independent of the other distinctions drawn so far, as the following Figure shows:

<i>Occurrent</i>	Instantaneous	unstructured	structured
Simple	×	×	×
Complex	×	×	×

⁵⁸ Later we argue that, in some sense, the particular stages of a gastroscopy may be said to instantiate the type 'gastroscopy'. Hence, we cannot say that the parts of a complex process do not instantiate the whole process. Rather, we call an occurrent complex if its stages, besides instantiating the whole occurrent type, *also* instantiate more specific types.

In the remainder of this chapter, we will focus on complex processes. Whereas all complex occurrents typically have several different parts, the parts of processes are also arranged in a typical order.

3. *Types and Instances of Occurrents*

We have appealed several times to the notion of *types* of occurrents. Types of occurrents differ from their instances in that they do not occur or happen. If something occurs, it is an instance, and if something is an instance of a type of occurrent, it occurs. There are several distinctions that are best drawn at the level of types. All particular instances of temporally extended occurrents take a determinate course in the sense that there are conditions for determining when they have begun and ended, and they all have an internal temporal structure. However, some occurrents have a *typical* internal temporal structure. When we specify the typical internal structure of an occurrent, we thereby characterize the type of which it is an instance. The typical course of a particular occurrent has the structure that it has by virtue of being an instance of a certain type. The type determines a structure insofar as there are standards according to which its instances may be judged to be typical or atypical, complete or incomplete. As we have already noted, it may be the case that a type is only instantiated by atypical or incomplete specimens. That a type of occurrent has a typical structure does not mean that its instances often or usually exhibit this structure.

3.1. *Types of Occurrents are not Instantiated by Continuants*

Types of occurrents are always instantiated by particular occurrents. This manner of speaking is not universally observed. It is sometimes said that types of occurrents are instantiated by the things which undergo, or participate in, such occurrents (see Rödl, 2005, 164). On this way of speaking, Socrates may be said to instantiate the type ‘going for a walk’ when he goes for a walk. This way of using the term ‘instantiates’ may have its origin in the fact that there can be no actual occurrent without there being something which undergoes, or participates in, this occurrent.⁵⁹ Hence, whenever an occurrent occurs, there will also be a continuant which it involves. However, this does not mean that it is the continuant

⁵⁹ With the possible exception of sounds; see Strawson, 1959, Chapter 2.

itself which instantiates the type of occurrent in question. Rather, types of occurrents are instantiated by individual *occurrents*, just as types of continuants are instantiated by individual continuants. Continuants participate in, or undergo, occurrents, but do not instantiate their types.

3.2. *How to Instantiate a Type of Occurrent*

Another way of confusing occurrents and continuants underlies a claim made by Rowland Stout to the effect that processes are really types of occurrents to which something happens when things undergo them (Stout, 1997). This is confused in several ways. First, what happens with the type when it is instantiated is, according to Stout, obviously itself an instance of a further type of occurrent. But this would mean that what happens between the continuant and the process must be a further process undergone by both of them, which in turn should be described as something to which something further happens. An infinite regress has blossomed. Second, as we saw earlier, types of occurrents are not temporal, even though their instances are: types are not instantiated in such a way that they occur in time. Entities which occur (occurrents) and entities which participate in this occurrence (continuants) are themselves never types but rather instances of types. Thus in order to understand what it takes to instantiate a type of occurrent, one needs to distinguish (at least) three sorts of beings: types of occurrents, their concrete instances, and the continuants which participate in these instances.

Continuants may undergo change, and when they do so, something happens. But this occurrent will not itself be something that changes. Granted, our language often permits us to say, e.g., that an activity becomes increasingly rewarding, or that it starts to become tedious. This however cannot mean that the activity in question would itself undergo a change. For, in this case, the change is really a change of the continuants which are participating in the activity (see Aristotle, *Physics* V, 2). That an activity is increasingly rewarding or tedious, for instance, simply means that the one who is engaging in it changes her attitude, or that the types of activities she is called upon to perform during one phase of an occurrent please her more or less than those types which she is called upon to perform during other phases.

4. *Complex Processes and Their Parts*

Processes are internally structured, temporally extended occurrents. *Complex* processes, by virtue of instantiating a certain type, have parts which are themselves extended occurrents, belong to different types, and occur in a certain typical order. For instance, knitting a wrist band is a complex process, since it may be broken down into temporally extended occurrents which are instances of different, more specific types. In order more precisely to distinguish complex from simple processes, we need to consider the way in which their types are specified. For this purpose it will be useful to consider two special cases of complex processes: *intentional actions* and *speech acts*.

4.1. *Recipes for Actions*

A simple and common way of specifying a type of complex action, such as knitting a wrist band or making an omelette, is to give a recipe. Recipes are structured to serve agents who possess certain basic capabilities and want to perform or at least initiate a complex process of a certain type. For this reason, a recipe will not explain every detail of what happens in the course of the process in question, but will only point out the sequence of steps that are basic relative to the normal agent; that is, the steps that a normal agent can immediately carry out without further instruction, preparation, or training. In this context, we may define a simple action as an action that need not be explained by a further recipe specifying its different components. For instance, it will be immediately clear what to do when told to move one's own hand. One might of course divide any such hand movement into further components, but this does not render the action of moving one's hand complex, since all its components are executions of the same basic capacity. If an agent knows how to move her hand from here to there, she will also know how to go on moving it. For this reason, the entire hand movement may be considered simple and not complex. Complex actions involve the actualization of different capacities at different times, and are specified by recipes; simple actions need not be specified by recipes, since they do not involve the actualization of multiple capacities (see Baier, 1972). Whether an action is simple or complex depends on the abilities of the agent. When we acquire basic abilities, such as speaking a language, it often happens that, through training, complex actions turn into simple actions. Initially, we may need detailed instructions

as to how to pronounce a certain word but, later, will be able to pronounce it without needing to reflect, step by step, upon these instructions.

What has been said about actions and recipes may be applied to processes in general. A process is simple if it does not involve further extended occurrences of different kinds. Thus, we may specify a recipe which includes all different simple processes that are typically involved in a complex process, and which determines their general order. This recipe will list the elementary steps that the complex process typically involves. It will often be clear in what sequence the steps are to occur, and how many times they may be repeated, but this need not always be the case.

4.2. *Regular Expressions*

In several respects, language use is paradigmatic for complex actions and processes. For some linguistic devices, such as words, sentences, and poems, there are explicit rules of typicality that determine, to some degree of precision, their internal structure. Simple processes in this context include utterances of syllables or writings down of letters of the alphabet. Accordingly, written documents may be compared to types of utterances. Like a type, a written text does not occur, but it specifies the structure of something that may occur: a sequence of elementary utterances. Without necessarily adopting the view that written texts simply *are* types of utterances, one may still say that they relate to utterances in a way similar to the way in which types of occurrences relate to their instances. Here, we are interested in only one of the similarities between texts and types.

In order to search a text file for the occurrence of a given word, one may write a computer program that parses all of a text's elementary constituents and checks whether they anywhere match a certain pattern. For instance, in order to search for the word 'bench', a program may check whether there is any point where the letters *b*, *e*, *n*, *c*, and *h* occur in that order. This program will identify a sequence of phonemes according to a rule that may also be used as a recipe for producing this sequence of phonemes.

Further, a more flexible program may be written in order to identify more general patterns. Such a program may for instance search for all words in a given text that begin with a 'b', contain three further letters of any kind, and then end with an 'h'. In order to give such search instructions in a compact and convenient way, so-called *regular expressions* have been

developed.⁶⁰ For example, one may adopt the convention that an expression like ‘b.{3}h’ shall represent any string that begins with a ‘b’, contains exactly three further letters of any kind and then ends with an ‘h’.

This procedure should admit of further generalization, so that complex processes of all kinds may be specified by using regular expressions. In order to write a program for identifying process occurrences by using regular expressions, a programmer would merely have to specify a set of variables and operators such as ‘.’ and ‘{3}’, and a set of constants that refer to simple occurrents. A generalized regular expression would thus match patterns of all kinds of occurrents, not only of linguistic utterances.

Such regular expressions will match actual processes not by describing the exact temporal duration and order of their stages. There are more appropriate and better known devices for describing and specifying a given process in such detail; examples include movies, calendars, or clocks. In contrast, recipes and generalized regular expressions only specify the *most general* order of steps that a complex process involves, and do not specify how much time the constituent occurrents will take. References to time may be added in a second step if necessary (think of ‘boil the egg for five minutes’ as part of a recipe).

The existing rules for constructing regular expressions are rather complex and will not be discussed in detail here. An example must suffice. At the beginning of a typical gastroscopy, a nurse will turn the patient onto her left side, a process which we will represent by the letter *a*. (For the sake of simplicity, we assume that being turned on one’s left side by a nurse is a simple process. Otherwise, it will be easy to supply detailed instructions on how it is done, and insert them in place of *a*.) Second, the endoscopist will usually spray the patient’s throat with a local anaesthetic (*b*), and in some cases she will apply a light sedative (*c*). During the insertion of the endoscope via the esophagus (*d*), the patient should swallow (*e*) several times until the endoscope has reached the bottom of the stomach (*f*). And so on. The initial segment of the operation can now be specified by giving the following generalized regular expression:

$$^abc?(de?)+?f$$

‘^’ marks the beginning of the process, telling us that *a* occurs at the very beginning of a typical gastroscopy. Immediately after *a*, *b* will occur

⁶⁰ For instance in the programming language PERL (see www.PERL.org).

exactly once. The question mark immediately following ‘*c*’ indicates that *c* may either be done once or may be omitted (i.e. *c* is optional). The parentheses group the expression ‘*de?*’ which refers to a complex subprocess consisting of the movement of the endoscope (*e*) and an optional swallowing (*d?*). This subprocess may be repeated any number of times, which is indicated by the ‘+?’ Here, the question mark has a different function than in the other two cases. It tells us that the process (*de?*) shall be repeated at most until *f* occurs; that is, until the bottom of the stomach is reached.

In this way, we may use regular expressions in order to specify the typical form of all kinds of processes. Since generalized regular expressions describe the typical course that a kind of occurrent takes, we may use them in order to specify standards of typicality that apply to these kinds of processes.

5. *Types: Their Parts and Their Instances*

Generalized regular expressions serve mainly to describe in outline the structure of a given type of process, and they apply only mediately to their instances. Only a *typical* gastroscopy includes all the steps in the exact way that is specified by the regular expression. It may well happen that the physician leaves out a step of the operation or calls it off before it is complete. When this happens, she will have been performing a gastroscopy, but she will not have done everything that belongs to a typical gastroscopy. As long as she was performing it, however, it was a gastroscopy that she was performing, since she will then have done everything that belonged to a gastroscopy up until that stage. Even if she has not yet completed the operation, she is already performing *it* and nothing less.

Against this, one might object that the physician does not perform a complete gastroscopy before she has actually done everything that belongs to an operation of that kind. However, this would be to confuse a type of occurrent with its instances. The type must include the steps, but its present instances need not. For when a physician performs a gastroscopy, she does not (indeed, cannot) immediately and simultaneously perform everything at once that belongs to such an operation, but must perform only one thing at a time. But this cannot mean that she is not in every moment, while performing the gastroscopy, performing all of the steps that a gastroscopy requires. The fact that the gastroscopy unfolds over time does not imply

that, before she has completed it, she is performing an *incomplete* gastroscopy. Even before it is complete, it is a gastroscopy that she is performing.

For if, in order to be said to be performing a gastroscopy, the physician would have to perform every step belonging to one, it would be impossible for anyone ever to *be performing* a gastroscopy – that is, to be *in the course of* performing one. For suppose for a moment that, as long as a physician has not yet performed every step belonging to a gastroscopy, she *is not yet performing* one. On the other hand, it should be clear that as soon as she has completed the last step, she *is no longer performing* a gastroscopy. This would mean that the only instant at which we could correctly say that she is performing a gastroscopy would be the instantaneous and infinitely short event that constitutes the end of the last step. But we certainly do not want to say that performing a gastroscopy takes no time. In order to prevent this unintuitive result, we will have to say that the physician already *is performing* the gastroscopy before she carries out the last step. This means that she also is already performing a gastroscopy even if she should ultimately fail to carry out the last step.

5.1. *Duration by Virtue of Type*

We see again that types of occurrents differ essentially from their instances. The specification of a type must be complete in order to be a specification of this type and not of another one. If we leave out the last step of a gastroscopy from a specification of this type of occurrent, we alter the specification in such a way that it turns into a specification of something else. In contrast, by leaving out one step of a *particular* gastroscopy, the occurrent in question does not turn into something else. For all along, up to the point at which the last step was left out, it was true to say that what was going on was a gastroscopy; this truth cannot suddenly turn into a falsehood. A token gastroscopy with a missing step is still a gastroscopy, but an atypical or incomplete one.

This means that it is not essential for a token occurrent that all parts belonging to its type actually be instantiated. The parts of a gastroscopy do not *immediately* belong to what the physician does when she inserts the endoscope, since the physician may do everything that is required *at this stage* without doing everything else that belongs to a gastroscopy. What she is doing right now has the other steps that are involved in a gastroscopy as its parts only because it instantiates a *type* to which these parts are

essential. Thus in some sense, what the physician does is only one single thing, moving the endoscope, whereas in another sense, she is performing a complex operation with several parts. What she does is complex only by virtue of its instantiating a complex type.

Does this not again imply that instances of occurrents are infinitely short? For if we may say that in one sense what the physician performs is only a simple hand movement, we may further say that in the same sense, what she is performing is only that part of this movement that happens exactly now, which is infinitely short.

However, since the very same movement instantiates both the type 'movement from here to there' *and* the type 'gastroscopy', it would be wrong to say that it is *only* a simple hand movement. Whereas it is correct that every bit of the movement, however short, will also instantiate the type 'gastroscopy', it is wrong to identify the duration of one of them with the duration of the entire operation, since in fact a gastroscopy typically takes 10 or 15 minutes. What the argument shows is rather that the duration of an event is strictly relative to the type that it instantiates. If we take something as an instance of 'movement from here to there', it may take two seconds, but if we take the very same occurrent as an instance of the type 'gastroscopy', it will take longer. It does not make sense to say that this token occurrent has its very own duration, independently of any type that it instantiates. Rather, occurrents have their duration only by virtue of being of a certain type.

Hence, the duration, as well as the internal temporal structure and components of a token occurrent, belong to this particular only by virtue of its being of a certain type. How long a given occurrent takes and what structure it has will depend on the type which we say that it instantiates. Taken as an instance of 'handling of the endoscope', what happens may last 3 minutes; taken as an instance of 'gastroscopy', it may last 13 minutes. This does not mean, however, that we may arbitrarily choose how to refer to something and thus arbitrarily determine how long it takes. We cannot invent the types that a particular instantiates; we can only choose among them. The movement of the physician, for instance, may be taken to instantiate either the type 'insertion of the endoscope', or the type 'gastroscopy', but it will never be correct to say that it instantiates a type such as 'having breakfast'. There is always a fixed range of types that an occurrent may be taken to instantiate, and which types make up this range depends on actual circumstances in the real world, not on our imagination or willpower.

5.2. *Past Occurrents*

We have claimed that a particular occurrent has its duration only by virtue of being of a certain type. Against this, it might be argued that there are past occurrents, which are tokens and have a concrete duration. Suppose that a physician in fact spends 20 minutes performing a particular gastroscopy. In this case, there will be a duration associated with the concrete instance, and it will not be the ten minutes that a typical gastroscopy takes. Hence, it seems that the duration of a token gastroscopy may differ from the duration that it has by virtue of its type.

But first, it is simply wrong that a past gastroscopy *has* a duration. The reason is that there *are* no past gastroscopies. By assumption, past gastroscopies are past and over, and thus one may only say that there *was* a past gastroscopy which *had* a certain duration. Past gastroscopies do not occur (but rather *have occurred*), and hence, there *are* no past instances of the type ‘gastroscopy’ (rather, there *were* such instances). That there are no past occurrents should be as obvious as, say, that future events are not the ones that have already happened.

To insist on the proper use of the past tense in this way may seem pedantic, but it has repercussions for the question whether a past, concrete gastroscopy may, as such, be said to have a concrete duration. For as long as the past gastroscopy was still occurring, it did not yet have its concrete duration. As long as it existed, it was not yet over, and something might have happened that would have made it longer. Shifting from present to past gastroscopies thus makes no difference, since in order to attribute a duration to a past gastroscopy, we have to situate ourselves in the past, as it were, and from that point of view we cannot yet know its duration. There can be no time at which there *is* a token gastroscopy that would, as this very token, have a concrete and fixed duration.

Second, it may well be that a gastroscopy in fact took 20 minutes. This will be an atypical duration for a gastroscopy, but that will not mean that it did not have this duration also by virtue of being of a type. It had this atypical duration only insofar as it was correctly taken to be an instance of the type ‘gastroscopy’, not insofar as it may also correctly have been taken to be an instance of the type ‘insertion of the endoscope’.

All in all, types of occurrents differ from their instances in the following important respects. First, the distinction between occurrents with and without an internal temporal structure is applicable strictly speaking not to tokens, but only to their types. Second, token occurrents may instantiate

several types at once, for instance the types ‘hand movement’, ‘insertion of the endoscope’, and ‘gastroscopy’. This means that, thirdly, token occurrents have their duration only by virtue of instantiating one of their several possible types. Taken as a simple hand movement, this token occurrent may be correctly said to take two seconds; but taken as a gastroscopy, the *very same* token occurrent may be taken to last for 10 minutes.

5.3. *Types of Occurrents Instantiated by Their Parts*

Put in a pointed way, this means that types of complex processes may also be instantiated by instances of their proper parts (see Allen, 2005, 23-37). Turning the patient on her side and applying the spray are only parts of a typical gastroscopy, but when they occur, they are also fully fledged instances of that type. For we may point at the event at any time and say: ‘what is occurring here and now is a gastroscopy’. If this is correct, it indicates an important difference between token occurrents and token continuants. For instance, the type ‘bench’ is instantiated only by the complete bench, not by any of its parts in isolation. Although a bench may lack parts, or be atypical and incomplete, yet still be what it is, it would not be correct that, say, one of its feet alone instantiates this type. In contrast, the very beginning of a gastroscopy, the nurse’s turning the patient on her left side, will already fully instantiate the type ‘gastroscopy’.

Alvin Goldman (1970) has claimed that this is not the case. According to Goldman, a token event does not instantiate several types at once, but rather, someone who performs a gastroscopy and during its course inserts an endoscope does two different things at once: first, she performs a gastroscopy, second, she inserts the endoscope. Now it is perfectly possible to do two things at once, for instance, to perform a gastroscopy while chatting with the nurse. But this is certainly not what happens when the endoscopist inserts the endoscope and performs the gastroscopy. Otherwise, the same logic will imply that when she inserts the endoscope thus far in the course of inserting it farther, she is doing two things at once. And this will quickly lead to the claim that everyone is always doing infinitely many things at once. A theory that leads to such a claim is surely neither useful nor representative of reality.

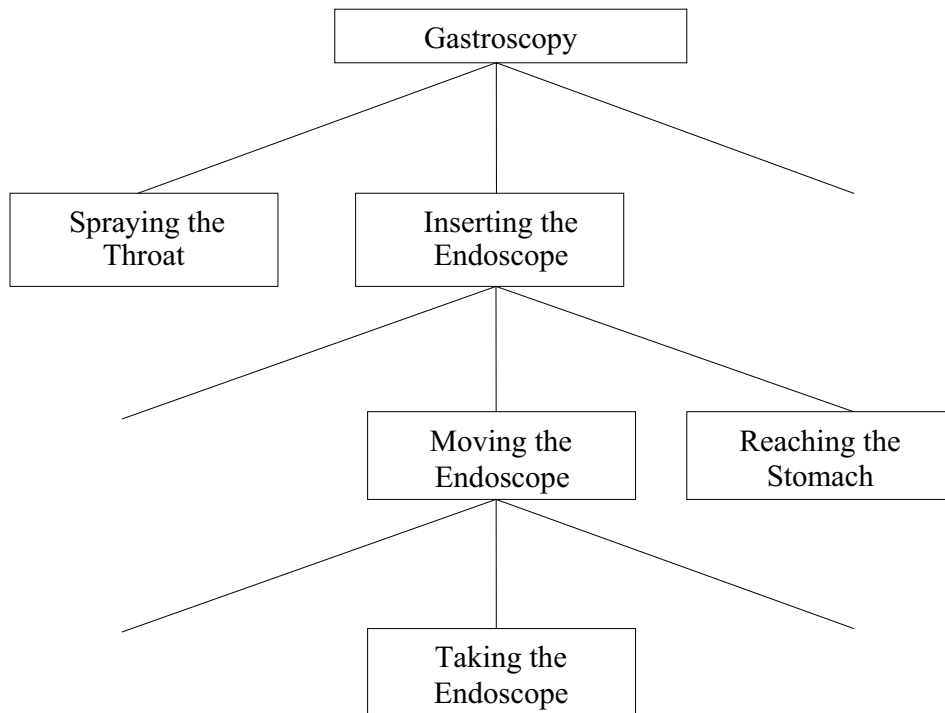
5.4. *A Tree Diagram*

The occurrent resulting when a doctor inserts an endoscope in order to perform a gastroscopy belongs to multiple types; to help us understand this we may think of the way in which a beaver may be said to belong to different types. A beaver is at the same time a rodent, a mammal, and a chordate, but this does not imply that a beaver is actually three *different* things at once. Rather, for this beaver, being a rodent is included in what it is to be a beaver. In a similar way, at some certain time instant, performing a gastroscopy is included in what it is, in the case of this endoscope insertion, to insert an endoscope, and we need not suppose that performing a single step and carrying out the whole procedure to which it belongs are two different simultaneous occurrents. Further, like the several types that a beaver instantiates, the several types that the movement of a doctor instantiates while she is performing a gastroscopy may be arranged in a tree diagram.

Types of complex processes may be specified by rules that describe the nature and order of the elementary steps that are involved. Since these steps are, during their occurrence, also instances of the respective parent process, we may take them to be leaf nodes that belong to the complex process as their root node in a tree diagram. Such a tree diagram will not have as much expressive power as a regular expression as defined above. We may indicate in a tree diagram which step involves what further steps, but not how often they shall be repeated, whether they are optional, or whether they are to be taken in an exact order.

Nonetheless, a tree diagram is a helpful tool for visualizing the structure of a complex process. Its root node will stand for the entire process, and it will branch into its immediate and possibly still complex parts. These parts may then have further leaf nodes. The ultimate leaf nodes will stand for elementary occurrents. As a rule, every token occurrent will also instantiate the types to which the respective parent nodes and their parent nodes refer.

The following, admittedly fragmentary and simplistic tree diagram, for instance, might represent a typical gastroscopy:



Whereas a regular expression only specifies the arrangement of the elementary steps that a complex process involves, the tree diagram bundles them into rough steps which are themselves subject to division into further, more basic steps. It looks a bit like the diagram in section 2.6 above, that is, like a taxonomy of different types of occurrents. On the other hand, it may seem to represent a partonomy; that is, a hierarchy of wholes and their parts. However, tree diagrams of the kind under consideration here neither represent taxonomies nor do they represent partonomies. Rather, they provide a visualization of the different ways of referring to the token occurrents that are typically involved in a complex process. Such a diagram does not tell us that every insertion of an endoscope is also a gastroscopy or vice versa; nor does it tell us that every insertion of an endoscope is part of a gastroscopy. All it says is that some particular token movement of an endoscope, if it occurs in the course of an endoscopy, may also be taken to be an instance of the types ‘inserting the endoscope’ and ‘gastroscopy’. What it represents is thus not a general hierarchy of types of occurrents, but only a hierarchy of types that one token may be said to instantiate at a time.

The diagram should not be taken to represent a partonomy because it is in any case problematic to speak of the *parts* of a token occurrent. As we have seen, particular instances of occurrents are complex only by virtue of instantiating a certain type. As a consequence, what parts a concrete token

occurrent may be said to have also depends on the type that it is taken to instantiate. A token occurrent may be said to have a part if its type has this part. However, types of occurrents do not have *temporal* parts, since they do not occur; only occurrents themselves have temporal parts. Nor is it clear in what other sense an abstract entity like a type may be said to have parts. The only thing that plainly has parts is the recipe by which we specify the type, but these parts are also not temporal parts. Further, as a matter of logic, nothing can be identical with its own proper part, since a proper part of X is defined as something *other than* X that is part of X. Neither the type nor the recipe is identical to their proper parts. But in the case of an ongoing gastroscopy we would want to say, first, that inserting an endoscope is part of, and something other than, performing the gastroscopy, since it may also occur in other contexts, but also that inserting this endoscope here and now is the very same as performing a gastroscopy here and now.

Thus, it seems that token occurrents may be identical to their proper parts. We should conclude that we had better avoid the word ‘part’ when it comes to instances of temporal entities. When we say that a token occurrent has parts only mediately, that is, by virtue of instantiating a certain type, we circumvent this difficulty, since the type is not a temporal entity and may thus be said to have non-temporal parts (which are determined by the parts of the corresponding recipe). In any case, the tree diagram should not be taken to represent a parthood, since its nodes stand for token occurrents, and these token occurrents should not be said to have parts (other than in the mediate sense, namely: by virtue of being of a certain type).

Let us note, in passing, that although our standard example of performing a gastroscopy is an intentional action, everything that we have said also applies at least to other natural processes, such as digestion, the movements of animals, and the growth of plants. Moving the hand is an instance of performing a gastroscopy not only because the physician *intends* to perform a gastroscopy when moving her hand. We may in fact say, without any conceptual difficulties, that an animal or plant, and even a machine, performs a complex task by performing steps that are involved in it, and draw a tree diagram representing the relations between the different ways of referring to what it does. For instance, what a hydrangea does when it grows may be divided into several elementary steps A, B, C, such that we may point at it and say that right now, what it is doing instantiates both the type A and the type ‘growing’.

5.5. *Necessary Incompleteness*

A token process must be taken to be already occurring before it is over, since it would not any longer occur when it is over. If the endoscopist inserts the endoscope at all, she is already doing it before she is done with it, since she would not be doing it any longer when the task is completed. We see what steps are involved in a complex process by looking at its type, and the specification of the type will tell us when this type of process is complete or over.

This is another respect in which temporal entities differ from non-temporal entities. Even if nothing should be perfect in this world, we may at least imagine a perfect thing, say, a perfect endoscope. Such an endoscope will possess all and only the features that an endoscope is supposed to possess. In contrast, it is simply impossible to imagine a perfect, that is complete, process. When we imagine a process, we must imagine it as going on, and as long we imagine it as going on, we will imagine it as not yet being over. As soon as it is over, it will no longer occur; and since for a process, to exist is to occur and to be complete is to be over, no complete process can possibly exist, as complete, at present. Hence, processes are present only as long as they are incomplete (see Aristotle, *Metaphysics IX*, 6, 1048b30). However, there are objections that may be raised against this view, which will be discussed in the following two sections.

5.6. *The Coast of Norway*

A token process appears to unfold in time just as, say, the coast of Norway extends in space. And although there is no single spatial location at which the coast of Norway is complete, the coast is nonetheless completely present throughout the entire extent of space that it occupies. Likewise, one may want to say that although a gastroscopy is complete at no single instant during its occurrence, it is complete throughout the entire stretch of time that it occupies.

This way of speaking seems to be plausible, but only on the basis of an illusion. We may see that the coast of Norway is completely present within a spatial region by traveling along this coast. Besides the three spatial dimensions that it occupies, there is a further dimension, time, which allows us to inspect all of its parts, not all at once, but one at a time. Further, that the coast of Norway is completely present throughout a spatial

region means that that there is a single instant at which it is everywhere in this region. That it is completely present means precisely that it is *now* here and there – that is, at one and the same time.⁶¹

In contrast, no token process is simultaneously at its beginning point, going on, and almost over, and there is no fifth dimension along which we could travel, as it were, in order to inspect all its parts. There is no ‘time’ at which all stages of a token occurrent could be simultaneously present.

5.7. *Four-dimensionalism*

Time is often taken to be a fourth dimension similar to the three spatial dimensions, and accordingly, occurrents are taken to be something like four-dimensional objects. This may do as a technique for mapping processes onto a four-dimensional coordinate system. Locations in Euclidean space may be represented by triples of numbers that describe them relative to the origin of a coordinate system, which will be represented by the triple (0,0,0). The location one unit left of the origin, for instance, may be represented by (-1,0,0). Three-dimensional objects can then be specified by sets of triples, such as $\{(0,0,0),(-1,0,0),\dots\}$. The same procedure may be applied in order to describe temporal entities. In order to do this, one may add a fourth number to each triple, representing a temporal instant relative to some temporal origin. Occurrents will thus be represented by sets of quadruples of numbers, referring to locations in a four-dimensional coordinate system.⁶²

One should not suppose, however, that nothing essential is lost when occurrents are represented in this way. What is lost when occurrents are transformed into sets of quadruples of numbers is precisely their *temporality*. Nothing about a quadruple of numbers in itself tells us which of the numbers refers to time, and without a convention according to which one of the numbers is to be read in a special way, the quadruple may as

⁶¹ Consider the fictional case of a coast segment that changes its shape and moves with us wherever we move. This coast segment would not be present throughout the spatial region in which we observe it, but it would first be here, then there. In order to know whether what we see is the complete coast of Norway, we must exclude this conceptual possibility.

⁶² This is not the same as to take snapshots of a complete situation containing three-dimensional objects and arranging them along a temporal axis, as it is done in Basic Formal Ontology. See *BFO*; Grenon, *et al.*, 2004; Grenon and Smith, 2004; Grenon, 2003. BFO does not identify continuants or occurrents with the sums of such snapshots.

well stand for spatial objects plus any other further dimension; say, their weight. Given that this is the case, consider what we are supposed to do when properly reading the fourth number. In order to read it, we have to relate the number to real or imagined temporal instants. We basically perform the same task as a DVD player: we map raw data, which is not an occurrent, onto a stretch of time so that an occurrent results. We map one of the four numbers onto real time, *first* reading one of the remaining triples, *then* the next one and so on. Since carrying out this very procedure is an occurrent, we re-introduce time in order to read the quadruple as a representation of something temporal. Note further that nothing about the set of quadruples of numbers in themselves tells us how *fast* we shall turn from one step to the next. The temporality of occurrents is thus in fact not preserved in their four-dimensional representation. Sets of quadruples only contain the data that must be mapped onto real time by performing a real process in order to represent something temporal. That is, although it may be very useful to isolate this data, quadruples of numbers should not be taken to represent all there is to occurrents. Occurrents, that is, are not four-dimensional entities.

If occurrents are not four-dimensional entities, however, there is no sense in which a process may be said to be completely present in the same way in which the coast of Norway may be said to be completely present. Rather, token processes are necessarily incomplete as long as they exist.

6. *Conclusions*

It has emerged that there are certain properties, such as structure and duration, which token occurrents have only mediately, that is, insofar as they instantiate certain types. Similarly, the distinction between instantaneous and extended occurrents can only be drawn properly at the level of types of occurrents. What happens right now when the physician performs a procedure is, under one of its possible descriptions, an instantaneous event; under other equally possible descriptions it is extended, elementary, or complex. Whenever we refer to some specific occurrent, we must refer to it as an instance of a certain type and, depending on this type, it will have different properties. For instance, referred to as a gastroscopy, what the doctor is doing right now will have a typical duration of 10 to 15 minutes, and it might turn out to have had an actual duration of, say, 13 minutes. Taken as an instance of inserting the endoscope, however, the very same occurrent will have a different typical

and actual duration. The different types which a token occurrent may instantiate may be brought into a system or hierarchy that may be represented by a tree diagram.

We have also shown that every currently occurring token process is necessarily incomplete, and that there *are* no complete processes. (There are occurrents other than processes, namely *energeiai*, that may be complete before they are over.) Past processes may be said to be complete, but only because they no longer exist. They *have* occurred in the past and *are* now complete. Further, past occurrents can only be understood as occurrents that once were present; that is, no reference to past occurrents will clarify the nature of present occurrents. Hence, that every past process has in fact had a specific structure and duration does not imply that present processes have such a structure and duration independently of their type.

Further, we argued that, although the concrete structure of token processes may be mapped onto a four-dimensional coordinate system, this should not be taken to imply that they are, in fact, four-dimensional entities. A set of quadruples of numbers counts as a representation of an occurrent only if there is a procedure by which the time index may again be mapped on real time instants. But first, this procedure will add back in the time that was lost in the representation, and second, the procedure itself is an occurrent. Hence, it is not possible that *all* occurrents should be encoded by sets of quadruples. Again, since occurrents are not four-dimensional entities, there is no sense in which they could be said to be complete as long as they exist.

Chapter 13: Bioinformatics and Biological Reality

Ingvar Johansson⁶³

Many bioinformaticians seem to shy away from believing that there is a mind-independent biological reality at all, or believing that we can have knowledge about such a reality. The aim of this chapter is to try to counteract this tendency, and it consists of two main parts. In the first part, I clarify three different positions in the philosophy of science with which it would be fruitful for bioinformaticians to become familiar. When they are spelled out in some detail, it becomes evident that these positions are mutually exclusive, but when seen only vaguely, the false impression may arise that one can sometimes rely on one position and sometimes on another. I label them *Myrdal's Biasism*, *Popper's Epistemological Realism*, and *Vaihinger's Fictionalism*, respectively, and I defend Popper's position.⁶⁴ In the second part of this chapter, I infuse new blood into the common semantic distinction between the *use* and *mention* of terms and concepts.⁶⁵ Both the red and the white blood corpuscles in this new fluid come from the philosophy of intentionality.

1. *Myrdal's Biasism*

Now and again we think of, and even perceive, the world in a way that is closer to how we would like it to be than how it really is. In such

⁶³ Slightly revised from the version originally published in *Journal of Biomedical Informatics* 39 (3:2006), 274-287; used with permission from Elsevier.

⁶⁴ The content of this paper has gradually come to fruition over the course of many conferences and workshops concerned with philosophy and informatics. The conference 'Ontology and Biomedical Informatics' in Rome (May 2005) finally triggered me to make these thoughts as clear as possible. Both biasism (but not Myrdal's) and Vaihinger's fictionalism were, quite independently of me, put on the agenda in Rome by Alexa McCray's talk 'Conceptualizing the World: Lessons from History'.

⁶⁵ I will deny my own preferences and use 'term' and 'concept' instead of 'word' and 'meaning', respectively, in order to conform to the usage of bioinformaticians. To a non-Platonist philosopher such as myself, the term 'concept' suggests too many allusions to entities that exist in some extratemporal realm of their own, independently of human beings. 'Meaning', on the other hand, has no such associations. Meanings exist directly only in people.

situations, we are biased. But how often does this occur? And what are the consequences of such bias for scientific research? One position in the philosophy of science can be captured by the following thesis and proposal:

- Thesis: Every conceptualization and theory is biased.
- Proposal: Admit that you are biased and make the causes of this bias (valuations, social positions and backgrounds, etc.) explicit, both to yourself and to your readers.

This position, nowadays widespread, was put forward already in the fifties by the economist Gunnar Myrdal (who shared the Nobel Prize in economics with Friedrich Hayek in 1974), but only as a thesis about conceptualizations in the social sciences (Myrdal, 1956, 1968, 1973 (Chapter 7)). Myrdal's views quickly reached the general philosophical audience thanks to Ernest Nagel's discussion – and criticism – of them in his classic *The Structure of Science*, 1961.

At the time Myrdal was writing, it was commonly assumed that scientists in their research activities ought to be, and mostly were, neutral with respect to valuations (Myrdal's term) and values that are not purely scientific. In criticism of this assumption, Myrdal claimed (a) that it is impossible for social scientists to free themselves from all such valuations, and (b) that valuations necessarily distort. According to Myrdal, since the value-neutral social scientist is a myth, social science is always more or less biased and more haunted by conflicts than the natural sciences are, and the only thing that scientists can do to become more objective is to find out and clarify, both for themselves and their readers, what kinds of valuations they bring to their research.

Looking at the historical development of the natural sciences, one can now add further that, even though there is much scientific consensus among natural scientists at most points in time, there is nonetheless a great divide between natural scientists belonging to different epochs (contrast Europe, for example, at the times of Newton and Einstein). Such differences, it can then be argued, are due not to the discovery of new facts but to the different cultural values of the centuries and scientists in question. In this way, many people have moved from Myrdal's own biasism, which is restricted to the social sciences, to the generalized version, which applies to all sciences that are not purely formal. Logic and

mathematics are mostly regarded as being outside the scope of biasism, but I have seen no claim that bioinformatics should be regarded as such.

As I will show, biasism (in whatever version) contains at least three serious philosophical flaws, each of which is sufficient reason to reject it. First, it makes no sense to speak of something being to the right if there is nothing that can be said to be to the left; similarly, it makes no sense to speak of bias if it cannot be contrasted with truth. In Myrdal's writings, biasism does not (and cannot without losing its sense) take the concept of truth wholly away. What it does do is to claim that we cannot *know* truths and that, therefore, we should speak of research results as being true-for-certain-valuations instead of being just true.

Biasism does not say that scientists are *sometimes* biased and so put forward distorted research results. The claim of biasism is that this is *always* the case; either only in the social sciences (the restricted thesis) or in all the non-formal sciences (the general thesis). Let me compare biasism in science with issues of legal jurisdiction. Judicial procedure seeks to find non-biased judges and jury members. If biasism were applied to such procedures, it would amount to the claim that there are no non-challengeable persons at all. Because of this generality, the thesis of biasism has to be applied to itself. It then implies the following disjunction: either biasism is false or it is true; but in the latter case it says of itself that it is biased and therefore false. That is, it is necessarily false. Therefore, of course, it should not be adhered to.

However, the self-referential paradox of biasism can be taken away. The defenders of biasism have merely to claim that their thesis lies outside the harmful influences of valuations and that they, therefore, are in a position to state one of the few known truths. Their thesis would then be: All theories are biased, except the theory of biasism. But now, another problem pops up. They have to explain why their thesis – a thesis which belongs to the sociology of knowledge is, in contradistinction to all other scientific and philosophical hypotheses, not influenced by valuations. If their thesis really is true, then it seems to be a mystery why this cannot also be the case with scientific assertions of other sorts. As far as I know, no one has solved this problem; I think it is unsolvable.

There are at least two reasons why many otherwise good researchers do not notice the paradoxical character of biasism. First, it seems to be natural for people who make assertions such as 'Humans are always fools' and 'Humans are always liars' to place themselves outside the scope of what they say; if not altogether, at least at the moment of making the assertion.

Those who have asserted ‘Humans are always biased’ might have followed this habit without noticing it. Second, in the case in hand, it is easy to make a so-called fallacy of composition. That is, from the fact that something is possible in *each* case, one falsely thinks that one can draw the conclusion that this something is possible in *all* cases taken collectively. Obviously, from the fact that in a marathon race *each* starting runner may win, one cannot validly draw the conclusion that *all* runners may win. Similarly, but less obviously, from the fact that *each* scientific hypothesis may be biased, one cannot by means of mere logic draw the conclusion that *all* scientific hypotheses may be biased.

A second argument against biasism is the following. The biasist proposal says that scientists should make the causes of their biases explicit; but according to the biasist thesis, even such a presentation of one’s bias must itself be biased and therefore false. Why? Because to state what has caused one’s bias is as much a hypothesis as are other empirical assertions. According to the thesis, it must be impossible to know what the true causes of one’s bias are. If biasism is true, researchers do not only automatically get a distorted view of what they study, they also get a distorted view of what has caused their distorted research results. There are, so to speak, distortions all the way down. Therefore, there is no reason to follow the proposal.

However, as in relation to the first flaw, the defenders of biasism may attempt to bypass this self-referential oddity by qualifying their position. Confronted by this second curious feature of their position, they may claim that their proposal makes good sense because researchers are *less* biased when they try to find truths about the causes of their bias, than they are when they try to find other scientific truths. For instance, it may rhetorically be asked: Isn’t it easier for an economist to find out what his sex, ethnicity, social background, and social valuations are than to find out how, in some respect, the market works? I have two counter-remarks to this rejoinder. First, the relevant problem is much harder than merely discovering facts about one’s social position and background. The real problem is to discover what causes distortions in one’s own research. In such an undertaking, one has also to take into account the fact that sometimes people with the same social position and background have different opinions. But second, and more decisive, is the fact that this qualification breaks the biasist frame. If there are degrees in the way researchers are biased, there are degrees of distortion in research results; and if there are degrees of distortion, there are degrees of being true or

false, i.e., degrees of being ‘truthlike’ (to anticipate the section on Popper’s realism below). But if there are such differences in degree of distortion between ordinary hypotheses and hypotheses about factors that cause bias, then there seems to be no reason why one should not be able also to detect such differences in truthlikeness between ordinary hypotheses.

The critique that I have presented so far takes seriously the fact that biasism puts forward an all-embracing thesis, which, in effect, replaces the notion of being true with the notion of being true-for-certain-valuations, which, in turn, can ground notions such as true-for-us and true-for-them. Such a replacement leads, as I have pointed out, to inconsistencies. This criticism, it has to be noted, by no means implies that we are never justified in talking about bias in science. In local cases, and having recourse to the notion of truth, we seem now and then to be justified in asserting that some scientists have been biased. But such local accusations of bias must be kept distinct from biasism, which contains a universal thesis.

For several decades now, biasism comes naturally to many people. One causal factor behind this fact might be the following. Nowadays, a large number of people in Western societies earn their living performing research or research-like activities in which the final research report takes the form of a consensus statement written by a group. This is true of public commissions of inquiry, be they initiated by the state or some regional or local authority; it is true of research departments in big firms; and of the managements of many research institutes. In such groups, after the research is performed, there comes a phase in which the final results are negotiated. This process can easily convey the false impression that there are no truths at all, only negotiations about truths and, therefore, only truths-for-certain-valuations, truths-for-us, truths-for-them. As far as I can see, many bioinformaticians have a similar kind of experience when they try to do justice to the advice and opinions of experts in various domains of knowledge.

Finally, let us for a while imagine that biasism has no self-referential problems. Nonetheless, another curiosity appears. All research needs a regulative idea, something that tells the researchers what to look for. Traditionally, the overarching regulative idea has been truth. This does not mean that truth has to be at the center in every phase and corner of research. For example, physicists may now and then be merely playing with possible solutions to some equations without, for the moment, bothering about *truth* at all. Similarly, some biologists may be merely playing with simulations of various biological processes; and researchers

in the humanities may be playing with certain possible interpretations of texts. This means only that there is, even in these playful situations, an indirect connection with the discovery of real natural laws or physical facts, the discovery of real biological processes, and the discovery of true interpretations, respectively. Today's science relies on a division of labor where not every part has to have a direct link to reality. This being noted, the third flaw of biasism can be stated as follows:

- Biasism wants science to get rid of the regulative idea of truth, but it has no adequate alternative to offer.

According to the biasist proposal, researchers should admit that they are biased and make the causes of their bias explicit. But what is the purpose of this proposal? Since a rational person should not seek truth if he or she firmly believes that one cannot even come closer to it because there is bias all the way down, the proposal in effect implies that researchers should exchange the regulative idea of truth for the idea of truth-relative-to-the-researcher's-valuations. The latter could be specified either as the idea that researchers should try only to promote their own long-term interests, or that they should, in the course of their research, only try to have as much fun as possible. Although such a substitution has no logical flaws, it amounts to a substitution of science with something else. It implies, contrary to Myrdal's intention – which was to promote objectivity – that researchers should be allowed to consciously deviate from data and to consciously ignore data that they suspect are problematic *because of their own valuations*. In secret, individual researchers may well have such goals, but these goals cannot possibly be made the public goal of science. Who, for instance, would fund researchers who say that they will use their research money only in order to promote their own egoistic interests or only in order to have fun?

Note that the remarks just made are not at all concerned with the question of how *research problems* within one's discipline *are chosen*. Such choices can, of course, be easily related to valuations and interests. This reminds us that several factors may account for the popularity of biasism, and that one is precisely a conflation of the choice of a regulative idea for one's research (truth, long-term subjective interest, short-term fun) with the choice of research problem. Another factor may be a neglect of the fact that all philosophically minded researchers have to face what might be called an existentialist choice: shall I primarily try to find the true solution

to my problem, or shall I primarily try to find a result that promotes my interests, or shall I primarily try to merely have fun? We can call this choice existentialist, since it is inevitably a personal choice that every researcher must make for herself. It can be made consciously, half consciously, or subconsciously, but it has to be made.

Yet another factor behind the popularity of the biasist proposal, perhaps the most influential, is the fact that there really is something to the proposal when it is viewed from the perspective of the *readers* of research reports. However, this requires that the proposal be put within a traditional truth-seeking framework – a fact which is not noted by the adherents of biasism. In what way can someone reading a research report be helped by coming to know the valuations of the researchers in question? On non-biasist premises, the answer is simple. As soon as there is a division of labor in the knowledge enterprise of a community, the sources of knowledge traditionally discussed in epistemology, namely reason and observation, are complemented by trust (in those who are providing information) (Coady, 1992; Kalman, 1999). In order for laymen to accept knowledge or information from researchers, and in order for researchers to accept knowledge or information from other researchers, the former have to trust the latter; in information science, knowledge engineers normally trust the domain experts. Therefore, the readers may be helped in this trust issue if each researcher states: ‘trust me or not; I have done my best to find the truth with ordinary methodologies, but if you suspect that I have distorted facts in order to further my interests, then I can tell you that my sex, ethnicity, social backgrounds, and social valuations are as follows: ...’ An example, for simplicity’s sake not taken from biology, will make the point more lucid.

Imagine the following situation. Two different investigations have been made about the income distribution for a certain kind of job. According to report A, the average income is 15% higher for men than that for women, but according to report B it is 25% higher. The researcher behind report A states that he is a male income statistician who thinks that men ought to have higher salaries than women, and that, in particular, a 15% difference is too little, whereas the person behind report B states that she is a female statistician who thinks that men and women ought to have the same salaries for the same kind of job, and that a 25% difference is far too much. Whose report should be trusted? Both reports cannot be true, although both can be false. In my opinion, if it is impossible to perform further independent investigations of this matter of one’s own, it is somewhat

rational to trust the person whose values (not sex) one shares. This being so, there is a kernel of truth in the biasist proposal that researchers should make their valuations, social positions and backgrounds, etc. visible, but this kernel has here been placed within a context where traditional truth seeking is taken for granted. That is, the researchers in the example have both asked themselves: ‘what is the truth, what are the facts?’, and their readers ask themselves: have the researchers really found the true income distribution? A researcher who suspects that he (or she) really unconsciously distorts facts ought to perform his (or her) investigation twice. He should first make it, so to speak, spontaneously, and he should then work through it once more with the conscious intention of trying to detect hitherto unconscious distortions.

2. Popper’s Epistemological Realism

Outside the philosophy of science, Karl Popper is most well known for his defense of democracy in *The Open Society and Its Enemies* (Popper, 1945). Within the philosophy of science he is best known for his falsifiability criterion and his advocacy of *fallibilism*. The former consists in the thesis that scientific hypotheses, but not metaphysical assertions, are falsifiable, and that, therefore, scientists (but not metaphysicians) are marked by the fact that they can state in advance what could make them regard their hypotheses as false. Fallibilism is the view that no presumed knowledge, not even scientific knowledge, is absolutely certain. In order to put his falsifiability criterion to real work, Popper connects it with some other general methodological rules. Here, however, I will present only his general epistemological realism.⁶⁶ Although I wholeheartedly accept this realism, I believe that his falsifiability criterion and its concomitant rules have to be rejected (Johansson, 1975). Thus, Popper’s general realism can be dissociated from his methodological rules, from his view that there is a gap between science and metaphysics, and that there is a criterion for detecting this gap. In particular, I will highlight a notion that is crucial to Popper’s realism. He verbalizes it in three different ways: truthlikeness, verisimilitude, and approximation to truth (Popper, 1972). I find this notion extremely important, but unduly neglected outside circles of Popper

⁶⁶ This realism is best spelled out in Popper, 1963, in particular in Chapters 1-4 and 10, and in Popper, 1972, in particular in Chapters 2, 5, and 7-9. His falsifiability criterion and most of his methodological rules are put forward in his 1959.

experts. The core of Popper's epistemological realism can be captured by the following thesis and proposal:⁶⁷

- Thesis: Every conceptualization and theory almost certainly contains some mismatch between theory and reality.
(Compare Myrdal: Every conceptualization and theory is biased.)
- Proposal: Seek truth but expect to find *truthlikeness*.
(Compare Myrdal: Make your valuations, social positions and backgrounds, etc. visible.)

Popper's epistemological realism combines fallibilism with the traditional idea that truth seeking has to be the regulative idea of science. The key to this mix is the notion of truthlikeness (verisimilitude, approximation to truth). The intuition behind this notion is easily captured. Consider the three assertions: (1) The sun is shining from a completely blue sky, (2) It is somewhat cloudy, and (3) It is raining; or at the assertions (1) There are four blood groups plus the Rh factor, (2) There are four blood groups, and (3) All blood has the same chemical composition. In either case, if the first assertion is true, then the second assertion has a higher degree of truthlikeness and approximates truth better than the third one. This is not to say that the second is *more likely to be wholly true* than that the third. The sentences 'X is *probably* true' and 'X has *probably* a high degree of truthlikeness' express relations between an assertion X and its evidence, whereas the sentences 'X is true' and 'X has a high degree of truthlikeness' express relations between the assertion X and facts (truthmakers) in the world. The former sentences express evidential relations, the latter express semantic-ontological relations;⁶⁸ the idea of truthlikeness belongs to a correspondence theory of truth.⁶⁹

⁶⁷ Of course, any epistemological realism presupposes a philosophical-ontological realism. With respect to the spatiotemporal world, Popper has a kind of level ontology (with which I wholly agree), according to which neither biological reality nor mental reality can be ontologically reduced to lower levels. Also, he thinks that thought contents have a kind of objective existence in what he calls the Third World (as contrasted with material reality, which makes up what he calls the First World and mental reality which forms the Second World).

⁶⁸ The possible conflation between being truthlike and being probably true comes more easily in some other languages. In German, for instance, the corresponding terms are 'wahrheitsähnlich' (truthlike), 'wahrscheinlich' (probable), 'Wahrheit' (truth), and 'wahr' (true).

At the end of better and better approximations to truth, there is of course truth. To introduce degrees of truthlikeness as a complement to the simple opposition between true and false is a bit – but only a bit – like switching from talking only about tall and short people to talking about the numerical or relative lengths of the same people. The difference is this. Length corresponds both to real comparative and numerical concepts, but there are no such concepts for verisimilitudes. All lengths can be linearly ordered (and thus give rise to a real comparative concept), and a general numerical distance measure can be constructed for them (which gives us a quantitative concept). Popper thought that such concepts and measures of degrees of truthlikeness could be constructed, but like many others I think that the ensuing discussion shows that this is impossible (Keuth, 2000). That is, we have only a qualitative or semi-comparative concept of truthlikeness. Some philosophers think that such a concept of truthlikeness can be of no use (*ibid.*, 198-9), but this is too rash a conclusion.

To demonstrate that even a semi-comparative concept of truthlikeness can be useful and important, I will use an analogy. We have no real comparative concept for geometrical shapes, to say nothing of a quantitative concept and measure. Nonetheless, we continue to use our qualitative concept of shape; we talk about shapes, point to shapes, and speak informally about similarities with respect to shape. Sometimes we make crude estimates of similarity with respect to shapes and are able on this basis to order a small number of shapes linearly (shape A is more like B than C, and A is more like shape C than D, etc.); we might be said to have a semi-comparative concept. In my opinion, such estimates and orderings of a small number of cases are also sufficient to ground talk of degrees of truthlikeness.

In the same way that a meter scale cannot be used before it has been calibrated in relation to something external to it, a standard meter, so the concept of truthlikeness of theories cannot be used until one has judged, for each domain in which one is working, some theory to be the most truthlike. In this judgment, the evidential relation stages a comeback. As I have said, truthlikeness informally measures the degree of correspondence with facts, not the degree of correspondence with evidence. Nonetheless,

⁶⁹ The correspondence theory of truth says that the truth of an assertion (truthbearer) consists in a relation to reality or in a correspondence with facts (truthmakers). Note that there can be no degrees of falsitylikeness; there are no non-existent facts with which an assertion can be compared. But, of course, one may use ‘being falsitylike’ as a metaphor for having a low degree of truthlikeness.

degrees of evidence must come into play when judging what shall be, so to speak, the standard meter for verisimilitude. Note that such judgments are commonplace decisions even for biasists and social constructivists. They are made every time some course book in some discipline is chosen to tell students some facts.

The notion of truthlikeness is epistemologically very important. Today's history of science tells us that it is no longer possible to believe that science progresses by adding one bit of truth to another in the way brick houses are built by laying bricks on top of each other. Whole theory edifices have often had to be revised, and new conceptualizations introduced; this sort of development will probably continue for a long time, perhaps forever. If, in this predicament, one has recourse only to the polar opposition between true and false, and is asked whether one believes that there are any true theories, be it in the history of science, in today's science, or in the science of tomorrow, then one has to answer: there are not. If, however, one has recourse to the notion of truthlikeness, then one can answer as follows.

There are, so far, no empirical theories known to be true in some absolute sense, but, on the other hand, there are not many absolutely false theories either. Most known theories in the history of science contain some degree of truthlikeness, even if only a very low degree. Today, however, some theories have what is probably a very high degree of truthlikeness. Why? Many modern inventions – and modern standardized therapies – are based on scientific theories, and it seems absurd to think that all such inventions in technology and medicine are based on theories with very low degrees of truthlikeness, to say nothing of the thought that these theories are mere fictions (see next section). Think, for instance, of travel to the moon, images from Pluto, computers, the internet, the GPS system, physiologic contraception, artificial insemination, and organ transplantations.

Let me now add a quotation from Popper in order to show how he himself summarizes his views on truthlikeness (1972, 335):

I have in these last sections merely sketched a programme of combining Tarski's theory of truth with his Calculus of Systems so as to obtain a concept of *verisimilitude* which allows us to speak, without fear of talking nonsense, of *theories which are better or worse approximations to truth*. I do not, of course, suggest that there can be a criterion for the applicability of this notion, any more than there is one for the notion of truth. But some of us (for example Einstein himself) sometimes wish to say such things as that we have reason to conjecture

that Einstein's theory of gravity is *not true*, but that it is a *better approximation to truth* than Newton's. To be able to say such things with a good conscience seems to me a major desideratum of the methodology of the natural sciences.

Just as in ethics there are people who only think in terms of white or black and who always want to avoid nuance and complication, so in science there are people who simply like to think only in terms of true or false. Not many decades ago scientists thought of their research only in terms of being certainly true; today, having familiarized themselves with the history of science, many – and especially in domains like informatics – think of it only in terms of being certainly false or certainly fictional (see next section). In neither of these positions – being certain that one has truth on one's side, or laying no claims to truth at all – must researchers fear criticism; but on fallibilist premises they must once again learn to do so.

Applying the notion of truthlikeness to the history and future of science allows us to think of scientific achievements in the way engineers think of technological achievements. If a machine functions badly, engineers should try to improve it or invent a new and better machine; if a scientific theory has many theoretical problems and empirical anomalies, scientists should try to modify it or create a new and more truthlike theory. As in engineering it is no sin to invent imperfect devices, so in science it is no sin to create theories that turn out not to be true. Rather, the sin in both cases is in not trying to improve on problematic machines and theories. Also, and for everybody, it is of course better to use existing technological devices than to wait for tomorrow's, and it is better to trust existing truthlike theories than to wait for the science of tomorrow.

Most rules have exceptions. Perhaps bioinformaticians, unlike scientists in other disciplines, need not bother about the history of science or think through the conflict between Popperian fallibilism and biasism? Isn't it enough for bioinformatics simply to systematize what the present-day experts in biology tell them? No, it is not. Biological knowledge grows rapidly, and even a young discipline like bioinformatics will no doubt soon have to revise some of its achievements in light of new biological knowledge. In the Gene Ontology, this is taking place before our eyes. For example, the constructors of GO list molecular functions as terms which have *obsoleted*.

3. *Vaihinger's Fictionalism*

In the 1920s and 1930s, Hans Vaihinger's book *The Philosophy of As-If* (1924) enjoyed much success. Viewed from one side this book speaks to the general positivist trend of those times; viewed from another side however it also fits well with the social constructivist trend of more recent decades. The essence of Vaihinger's position is:

- Thesis: Absolute truth, if such there is, is not attainable.
(Compare Popper: There is absolute truth, but it is probably not attainable.)
- Proposal: Regard your theories as referring to fictions; don't concern yourself with truth and falsehood.
(Compare Popper: Regard your empirical theories as referring to the world; try to find out if they are false.)

Vaihinger holds that there is only one kind of real entity, the contents of our sensations (this is the positivist side of his thinking). Things and persons in the ordinary sense, matter and energy as spoken of in physics, and things in themselves as postulated by some philosophers, are all merely fictions. Nonetheless, there are reasons to live *as if* many of such entities are real; the expression 'live as-if X exists' at the heart of Vaihinger's philosophy should be understood as follows:

- If there were Xs and we knew it, then we would have to expect some specific things to happen, and, also, we would have to act in some specific ways. In fact, however, we know that there are no Xs. Nonetheless we ought to create expectations and act *as if* there were Xs.

In some parts of his book, Vaihinger makes clear distinctions between (i) hypotheses (which are directed towards reality and demand verification), (ii) semi-fictions (which abstract away some known features of an entity, as for example the irrationality of humans is abstracted away in the concept of *homo economicus*), and (iii) pure fictions (which are based on no abstraction of this sort); but in the end he turns everything (except the contents of sensations) into pure fictions and says (1924, 108):

we are able ultimately to demonstrate that what we generally call truth, namely a conceptual world coinciding with the external world, *is merely the most expedient*

error. [...] So-called agreement with reality must finally be abandoned as a criterion.

He stresses the importance and necessity of postulating fictions in all areas of life, practical, scientific, as well as ethical. Like contemporary Anglo-American social constructivists, he implicitly takes for granted that we can communicate with each other about such fictions, i.e., he implicitly regards communication as real.⁷⁰ Since the contents of sensations play a subordinate role in his philosophy, it is no accident that his ideas can be summarized in such a way that they become, as here, lumped together with those of present-day social constructivists.

It is interesting to note how similar Vaihinger's and Popper's theses are and, despite this, how dissimilar their proposals are. In my opinion, the small difference between their theses is of no importance at all. Even if Vaihinger had subscribed to the view that there is some low probability that absolute truth is attainable, I am sure that he would have put forward the same fictionalist proposal. Conversely for Popper, even if he had thought that absolute truth is in principle unattainable, he would still have put forward the same falsificationist proposal. What, then, makes Vaihinger and Popper differ so radically in their proposals? My answer is, in short: Vaihinger's lack of the notion of truthlikeness.

False and fictional assertions are in one respect different and in another similar. They are different in that it is possible to tell a lie with a false assertion but not with a fictional one. When we are lying we are presenting as true an assertion that is false, but fictional assertions are beyond the ordinary true-false dimension. The two are similar in that neither refers to anything in reality that corresponds exactly to the assertion in question. A false empirical assertion lacks a truthmaker as a matter of fact; a fictional assertion cannot possibly have one. Therefore, it is easy to confuse the view that all theories are false with the view that all theories are about fictions. Nonetheless, it is astonishing how easily Vaihinger goes from falsehood to fictions. Why does he not believe that there can be degrees of fictionality? The less that has been abstracted away in a semi-fiction, the closer an assertion about it is to a hypotheses, and the more that has been abstracted away, the closer an assertion about it is to a purely fictional assertion. Assertions about semi-fictions might be said to be semi-true, and

⁷⁰ It should be noted that some French post-structuralists, e.g., Derrida, even regard the idea of communication as a fictional idea, and they communicate this thesis in many books.

since being semi-true takes degrees, we have hereby simply created another name for truthlikeness.

If one assertion is more truthlike than another, then it is by definition also less false. However, this falsity content (to take an expression from Popper) can easily be turned into fictionality content, whereupon a more truthlike assertion can also be said to be a less fictional assertion. When we are reading about, say, Sherlock Holmes, we have no difficulty in placing this fictional character in a real setting, London between 1881 and 1904. Not everything is fictional in many works of fiction, and we often have no difficulty in apprehending mixtures of real and fictional reference. Something similar is true when one reads about the history of science. For example, when I read about the false hypothesis that there is a planet Vulcan between Mercury and the Sun, which might explain some seeming falsifications of Newtonian mechanics, then I had no problem in taking Vulcan to be a fictional entity postulated as existing in the real solar system in about the same way as I take Holmes to be a fictional character in a real London. When I read about the false hypothesis that there is a chemical substance, phlogiston, which is exuded by things that are burning (where in truth, as we now know, oxygen interacts with the things in question), then I have no problem in taking phlogiston to be a putative fictional substance in the world of real burnings. When I read about Galen's view that (what we call) the arterial system contains *pneuma* or *spiritus*, then I have no problem in taking this *pneuma* to be fictional, but the arterial system to be real. Those who write about the history of science often make the reader look upon statements which were once false assertions as being assertions about fictions. In retrospect, we should look upon superseded theories as having mixed reality and fiction in something like the way reality and fiction can be mixed in novels. This is to give fictions their due place in science, but such local uses of fictions must be kept distinct from fictionalism, which contains a universal thesis.

I will end this section on fictionalism with the kind of remark with which I started my criticism of biasism. Apart from all other curiosities, fictionalism is self-referentially inconsistent. Fictions are created, but if everything apart from the contents of our sensations is a fiction, then there is nothing except such content that can create the fictions. However, contents of sensation do not have such a capacity. Unfortunately, Vaihinger and most fictionalists do not see the need for this kind of self-reflection.

4. *Use and Mention: In the Light of an Optical Metaphor*

When we look at things such as stones, trees, and walls, we cannot see what is on the other side. But things like water and glass are such that we can look through them to the other side. In the case of spectacles, microscopes, and telescopes, this feature is extremely useful. By *looking through* their lenses, we are able better to *look at* something else. This being-aware-of-*x*-through-*y* phenomenon is not restricted to the visual sense. It can be found in the tactile realm as well. One can grip a tool and feel the tool against the palm, but when one is very good at using such a tool, this feeling disappears. Instead, one is primarily aware of that in the world which the tool affects. For instance, when one is painting, say, a wall with a brush, one is only (if at all) indirectly aware of one's grip on the brush, and is primarily aware only of what one is painting. One is *feeling* the surface of the wall *through* the brush. What glasses are for people with bad sight, the white cane is for blind people.

Speech acts, acts of listening, writing, and reading acts – in short, language acts – are, like glasses and the white cane, tools for improving our everyday life. They can be used to convey and receive information, to give and take orders, to make and apprehend emotional outbursts, and to do very many other things. Even though language acts do not have the same robust material character as tools have, they nonetheless display the same feature of being able to be both looked at and looked through. In the former case, one is directly aware of the linguistic entities as linguistic entities, but in the latter case one is aware of them at most indirectly. When, for example, one is conveying or receiving information in a language in which one is able to make and understand language acts spontaneously, one is not looking at the terms, concepts, statements, and propositions in question; nor is one looking at grammar and dialects. Rather, one looks through these linguistic entities in order to see the information (facts, reality, or objects) in question. We are looking at linguistic entities, in contrast, when for example we are reading dictionaries or examining terminologies. If I say: 'look, the cat has fallen asleep', I want my use of the term 'cat' to be transparent and to help the person I am addressing to get information about a state of affairs in the world. But if I say 'In WordNet, the noun 'cat' has 8 senses', then I want someone to look at the term 'cat'.

My distinction between looking *at* and looking *through* is similar to the traditional distinction between the *use* and *mention* of linguistic entities,⁷¹ and it applies both to factual talk and to reading novels. In fictional discourse, terms are *used* as much as they are in talk about real things, but they are used in a very special way. Fictional discourse is *about* fictional characters; it is not about terms and concepts. In fact, we are typically using the same terms and concepts in both fictional and factual discourse.

When one is not using lenses, one can look at them and investigate them as material objects in the world. One can, for instance, try to find out what their physical properties and internal structures are like. In the world of practice, one makes such investigations of tools only when they are not functioning properly and are in need of repairing. Something similar holds true of terms and concepts. Only when our language acts are not functioning well – think for instance of learning a new language – do we normally bother to look *at* terms and concepts in dictionaries.

Furthermore, we are able to switch quickly between looking through and looking at things. Car drivers should look through, not at, the windshield, but when driving they should also have the ability to take a very fast look *at* their windshield in order to see whether, for instance, it has been damaged by a stone. Something similar is true of people using a foreign-language dictionary. They should be able to take a look at a certain foreign term and then immediately start to look through it by using it. Let me summarize:

1. In the same way that we can both look at and look through many material things, we can both look at and look through many linguistic entities.
2. In the same way that we can quickly switch between looking at and looking through glass, we can also quickly switch between looking at and looking through linguistic entities.

And let me then continue the analogy by adding still another similarity:

3. In the same way that consciously invented material devices for being-aware-of-*x*-through-*y*, such as microscopes and telescopes, have provided new information about the world, so consciously

⁷¹ I do not regard the distinction between use and mention as the same distinction as that between object language and meta-language. The use-mention distinction does not split ordinary language into distinct levels.

invented linguistic devices for being-aware-of- x -through- y , such as scientific concepts, have provided new information about the world.

By means of the *invention* of new concepts, we can sometimes *discover* hitherto completely unnoticed facts. Often, we (rightly) regard discoveries and inventions as wholly distinct affairs. Some things, such as stones, can only be discovered, not invented; others, such as bicycles, seem only to be inventions. One person might invent and build a new kind of house, and other persons may later discover it; but the first person cannot both invent and discover it. These differences between inventing and discovering notwithstanding, devices for being-aware-of- x -through- y present an intimate connection between invention and discovery. By means of new inventions of the being-aware-of- x -through- y type, we can discover x . There are many x 's that we can discover only in this way.

The third point above should be understood partly in terms of the notion of truthlikeness: if an existing conceptual system is faced with a conflicting conceptual system which has a higher degree of truthlikeness, the latter should supersede the former. But, conversely, the notion of truthlikeness should also be understood by means of the distinction between looking at and looking through. I introduced the idea of truthlikeness with the three assertions 'The sun is shining from a completely blue sky', 'It is somewhat cloudy', 'It is raining', and I said that, *given that the first assertion is true*, the second one seems intuitively to be more truthlike than the third. A standard objection to such a thesis is that this sort of comparison can show us nothing relevant for a correspondence theory of truth, since what we are comparing are merely linguistic entities (assertions). However, this objection overlooks the distinction between looking at and looking through. Looking at the assertions allows us to see only similarity relations between the assertions themselves; but when we have learned to switch from looking at to looking through such assertions – at the reality beyond – then we can coherently claim that the second is more truthlike than the third.

In the same way that our choice of lens may determine what we are able to see, so too, our choice of concepts determines what we can grasp. However, this is no objection to the thesis of epistemological realism to the effect that we have knowledge about the world: it does not render truth a social construction. When, through a concept, we look at and grasp something in the world, this concept often (i) *selects* an aspect of the world, (ii) *selects* a granularity level (for instance, microscopic or

macroscopic), and (iii) *creates* boundaries where there are no pre-given natural boundaries. The concept nonetheless (iv) *does not create* this aspect, this granularity level, or what is bounded. Think of the concept *heart*. It selects a biological aspect of the human body; it selects a macroscopic granularity level; and it creates a boundary line between the heart and its surroundings which does not everywhere track physical discontinuities as for example where the heart meets the aorta and the veins (Smith, 2001). But, nonetheless, our invention of the concept *heart* does not create our hearts, and there were hearts many millions of years before there were concepts.

5. *The Fallacy of Mixing Use and Mention*

All ontologies in information science contain terms. The builders of such ontologies look mainly *at* the terms in question, whereas the users of ontologies look mainly *through* them. Like the users, the experts in the various specialized domains of knowledge generally look through the terms. However, an ontology such as WordNet presents a special case, for (if it is to be called an ontology at all) it is an ontology of terms and meaning; it is like a dictionary, not like a taxonomical textbook. In its treatment of the term ‘cat’, WordNet begins as follows:

The **noun** ‘cat’ has 8 senses in WordNet.

1. **cat**, true cat – (feline mammal usually having thick soft fur and being unable to roar; domestic cats; wildcats)
2. **guy**, **cat**, **hombre**, **bozo** – (an informal term for a youth or man; ‘a nice guy’; ‘the guy’s only doing it for some doll’) (*WordNet*)

It is doubly clear that the term ‘cat’ is mentioned and not used in WordNet. Both the scare quotes around the term ‘cat’ and the fact that it is preceded by the term ‘noun’ make it clear that WordNet contains no talk of real cats; both scare quotes and context are able to disambiguate between use and mention. Here, therefore, matters are clear. In many biomedical ontologies, however, use and mention are systematically confused.

The Gene Ontology Consortium asserts that ‘[t]he Goal of the Consortium is to produce a structured, precisely defined, common, controlled vocabulary *for describing* [italics added] the roles of genes and gene products in any organism’ (Gene Ontology Consortium, 2000). That is, it is not an ontology for looking *at* terms but for looking through terms.

GO consists of three different ontologies, one for cellular components, one for molecular functions, and one for biological processes. One graph in the latter ontology (as it looks when this was originally written in June 2005 – the problem has since been corrected) can be reproduced as in the figure below; it is to be read from bottom to top. The original graph contains arrows representing the subsumption (*is_a*) and part-whole relations (*part_of*):

```

Gene_Ontology
  part_of
    biological process
      is_a
        physiological process
          is_a
            metabolism
              is_a
                nucleobase, nucleoside, nucleotide, and nucleic acid metabolism
                  is_a
                    transcription
                      is_a
                        transcription, DNA-dependent
                          part_of
                            transcription initiation (GO: 0006352)

```

When a user of the GO reads this he is, I am sure, looking through the terms. That is, he reads it as signifying something like: ‘each transcription initiation is part of a DNA-dependent transcription, which is a kind of transcription, which is a special kind (nucleobase, etc.) of metabolism, which, like all metabolisms, is a physiological and biological process’. So far so good, but I have stopped at ‘biological process’. What about the last step? Reading it in the same way would yield: ‘Each biological process is part of the Gene Ontology’. But this is obviously false. It should instead be read: ‘the *term* biological process’ is part of the Gene Ontology’s hierarchy of *terms*. Thus use and mention of ‘biological process’ are here mixed. When one reads the ontology from the bottom up and arrives at ‘biological process’, this term should be regarded as *used*, but when one continues reading upwards, it should be regarded as *mentioned*.

Since, as I said earlier, most people are able to switch between looking through and looking at terms, the human users of the GO may perhaps

without noticing do so also here, and no harm is done. However, automated information-extracting systems are not able to make such switches. Obviously, GO would be a better construction without this mixture of use and mention. As the graph stands, it allows a fallacious inference to the effect that, if something is a biological process then it is part of a certain human artifact called the Gene Ontology. This might be called the fallacy of mixing use and mention. The solution is easy: let the graph end in ‘biological process’.

The same kind of fallacy appears as well (at least in June 2005 – this problem has still not been fixed) in the CRISP (Computer Retrieval of Information on Scientific Projects Thesaurus).⁷² There, one finds subsumption relations which can be represented as in the hierarchy below (to be read from the bottom upwards):

```

immunology
  is_a
  antigen
    is_a
    allergen
      is_a
      airborne allergen
        is_a
        pollen

```

Here, ‘antigen’ should be *used* in relation to ‘allergen’ (‘Each allergen is an antigen’), but *mentioned* in relation to ‘immunology’ (‘The *term* antigen is an immunological *term*’). ‘Allergen’ is a term among other terms in the field of immunology, whereas allergens themselves are among the entities that immunology studies.

The Health Level 7 Reference Information Model (HL7 RIM), also, conflates use and mention, with the unfortunate result that the users of the RIM are told by its authors that the RIM *cuts them off from the world*:

Act as statements or speech-acts are the only representation of real world facts or processes in the HL7 RIM. *The truth about the real world is constructed through a combination (and arbitration) of such attributed statements only, and there is no class in the RIM whose objects represent ‘objective state of affairs’ or ‘real processes’ independent from attributed statements. As such, there is no distinction*

⁷² <http://crisp.cit.nih.gov/Thesaurus/index.htm>. Accessed February 4, 2008.

between an activity and its documentation [italics added]. Every Act includes both to varying degrees. For example, a factual statement made about recent (but past) activities, authored (and signed) by the performer of such activities, is commonly known as a procedure report or original documentations (e.g., surgical procedure report, clinic note etc.) (HL7 RIM).⁷³

6. Use and Mention: In the Light of a Good Philosophy of Intentionality

As I pointed out in section 2.1, both terms (words) and concepts (meanings), as they are most commonly used, are invisible, since we most often look *through* them at the entities to which they refer. The need for a distinction between term and concept arises as soon as we discover a synonymy, be it between two terms in the same language or in different languages. For we then have to specify what makes the terms differ and what makes them similar in meaning (i.e., synonymous). The terms differ because they are constituted by different syntactic unities such as letters or words conceived of as purely graphical or acoustic patterns, and they are synonymous (as we say) because they express the same concept. A term is a fusion of a syntactic unity and a concept.⁷⁴ One looks *through* the concept, not through the syntactic unity, i.e., concepts are to terms what lenses are to glasses, microscopes, and telescopes.

The optical metaphor of looking through concepts is sustained by a certain approach in the philosophy of intentionality. The term ‘intentionality’ was introduced into contemporary philosophy by Franz Brentano in the nineteenth century. It refers to phenomena such as perceiving, thinking, reading, and desiring. Intentional phenomena have in common the feature that they contain a *directedness* towards something. Mostly, it is a directedness that originates in a person who is in a so-called intentional state, or who performs an intentional act, towards something else. There are, however, different opinions on how to analyze intentional

⁷³ The quoted statement and others in the documentation of the HL7 RIM are criticized in Vizenor, 2004.

⁷⁴ Those who are amenable to Ferdinand de Saussure’s linguistics can read the last sentence as follows: A sign is a fusion of a signifier and what is signified. Let me add that Saussure consciously abstracted away from his studies all questions concerning looking-through at referents. Some of his present-day followers, however, seem to take the position (criticized in this paper) that there simply are no referents.

phenomena. In my opinion, Edmund Husserl (1970) and John Searle (1983) have come the closest to the truth.⁷⁵ Let me quote Searle (18-19):

it is at least misleading, if not simply a mistake, to say that a belief, for example, is a two-term *relation* between a believer and a proposition. An analogous mistake would be to say that a statement is a two-term relation between a speaker and a proposition. One should say rather that a proposition is not the *object* of a statement or belief but rather its *content*. The content of the statement or belief that de Gaulle was French is the proposition that the de Gaulle was French. But that proposition is not what the statement or belief is about or is directed at. No, the statement or belief is about de Gaulle ...

Intentional phenomena are marked by a tripartition between (intentional) act, (intentional) content, and (intentional) object. Assume that you are reading a physician's report about your heart, which tells you that your heart has some specific features. At a particular moment, there is then your reading *act* and what you are reading about, the intentional *object*, i.e., your heart and its properties. But since your heart exists outside of your reading act, there must be something within the act itself in virtue of which you are directed towards your heart and its properties. This something is called the *content*; in assertions, it consists of propositions.

According to many non-Husserlian and non-Searlean analyses of intentionality, you are in your act of reading directed only towards the proposition, but then there is outside your awareness also a relation of representation between the proposition (the content) and the object (your heart). According to Husserl and Searle, on the other hand, you are, while, reading directed towards your heart (object) by means of the proposition (content). The first kind of analysis leaves no room for any sensible talk of looking through concepts and propositions, but Husserl's and Searle's analyses do. Though Husserl's and Searle's theoretical frameworks differ in other respects, both of them make it reasonable to believe that the metaphorical distinction between looking at and looking through concepts can be embedded within a truly theoretical framework.

⁷⁵ In this respect see Searle 1983, p. 18-9, 57-61, 97, and Husserl 1970, Investigation V, §11 and the appendix to §21 ('Critique of the 'image-theory' and of the doctrine of the 'immanent' objects of acts'). Despite later changes of opinion, Husserl retains his belief in the feature of intentionality that I will highlight.

7. *The Last Word and the Last Word but One*

In the first part of this paper I advocated Popper's realism, in particular his notion of truthlikeness. In the second part I advocated a Husserl-Searlean analysis of intentionality, in particular the view that in assertions one is directed towards the world by looking through terms and concepts. Now, in order to forestall the possible criticism that I cannot explain and make sense of the position from which I am talking, I want to bring in another prominent thinker, Thomas Nagel. I regard myself as speaking from the kind of naturalist, rationalist position that Nagel has tried to work out in his *The View from Nowhere* (1986) and *The Last Word* (1997). Below are two quotations. The first is from the introduction to the latter book, and the second is its concluding paragraph.

The relativistic qualifier – 'for me' or 'for us' – has become almost a reflex, and with some vaguely philosophical support, it is often generalized into an interpretation of most deep disagreements of belief or method as due to different frames of reference, forms of thought or practice, or forms of life, between which there is no objective way of judging but only a contest for power. (The idea that everything is 'constructed' belongs to the same family.) Since all justifications come to an end with what the people who accept them find acceptable and not in need of further justification, no conclusion, it is thought, can claim validity beyond the community whose acceptance validates it.

The idea of reason, by contrast, refers to nonlocal and nonrelative methods of justification – methods that distinguish universally legitimate from illegitimate inferences and that aim at reaching the truth in a nonrelative sense. Those methods may fail, but that is their aim, and rational justification, even if they come to an end somewhere, cannot end with the qualifier 'for me' if they are to make that claim (1997, 4-5).

Once we enter the world for our temporary stay in it, there is no alternative but to try to decide what to believe and how to live, and the only way to do that is by trying to decide what is the case and what is right. Even if we distance ourselves from some of our thoughts and impulses, and regard them from the outside, the process of trying to place ourselves in the world leads eventually to thoughts that we cannot think of as merely 'ours'. If we think at all, we must think of ourselves, individually and collectively, as submitting to the order of reasons rather than creating it (*ibid.*, 143).

Reason, Nagel says, has to have the last word. However, this statement needs to be qualified. As the logician Per Lindström notes with regard to Nagel's book: 'reason has the last word – or perhaps only the last but one, since reality, reason tells us, has always the absolutely last word'

(2001, 3-6). Not only for biologists, but even for bioinformaticians, biological reality has the last word; notwithstanding the fact that bioinformaticians need not consult it too often. Mostly, they can trust the domain experts who provide them with their information.

References

- Aertsen, J. A., et al. (1998). Transzendental, Transzendentalphilosophie. *Historisches Wörterbuch der Philosophie, Vol. 10*, 1358-1465. Basel: Schwabe Verlag.
- Allen, R. (2005). The Mereology of Events. *Sorites*, 16, 23-37.
- Angelelli, I. (1967). *Studies on Gottlob Frege and Traditional Philosophy*. Dordrecht: Reidel.
- Anscombe, E. (1957). *Intention*. Oxford: Basil Blackwell.
- Anselm of Canterbury (1938-1961). *Opera omnia ad fides codicum recensuit*. (F. S. Schmitt, Ed.). Edinburgh: Thomas Nelson & Sons.
- Aquinas, T. (1964). *In duodecim libros Metaphysicorum Aristotelis exposition*. (M.-R. Cathala & R. M. Spiazzi, Eds.). Turin/Rom: Marietti.
- Aquinas, T. (1965). *In octo libros Physicorum Aristotelis expositio*. (M. Maggìolo, Ed.). Rome: Marietti.
- Aristotle (1936). *Aristotle's Physics*. (W. D. Ross, Ed.). Oxford: Oxford University Press.
- Aristotle (1958). *Aristotle's Metaphysics*. (W. D. Ross, Ed.). Oxford: Oxford University Press.
- Aristotle (1984). Categories. *The Complete Works of Aristotle* (Revised Translation, J. Barnes, Ed.). Princeton: Princeton University Press.
- Aristotle (1984a). Metaphysics. *The Complete Works of Aristotle* (Revised Translation, J. Barnes, Ed.). Princeton: Princeton University Press.
- Aristotle (1984b). Rhetorics. *The Complete Works of Aristotle* (Revised Translation, J. Barnes, Ed.). Princeton: Princeton University Press.
- Aristotle (2001). *On The Parts of Animals I-IV*. (J. G. Lennox, Trans.). Oxford: Clarendon Press.
- Armstrong, D. M. (1978). *Universals and Scientific Realism. Vol. I: Nominalism and Realism*. Cambridge: Cambridge University Press.
- Armstrong, D. M. (1978b). *Universals and Scientific Realism. Vol. II: A Theory of Universals*. Cambridge: Cambridge University Press.
- Armstrong, D. M. (1997). *A World of States of Affairs*. Cambridge: Cambridge University Press.
- Asplund, J. (1968). *Sociala egenskapsrymder*. Uppsala: Argos.
- Baader, F., et al. (Eds.). (2003). *The Description Logic Handbook*. Cambridge: Cambridge University Press.

- Baier, A. C. (2003). Ways and Means. *Canadian Journal of Philosophy*, 1, 275-293.
- Barker, R. G. (1968). *Ecological Psychology: Concepts and Methods for Studying the Environment of Human Behavior*. Stanford: Stanford University Press.
- Beard, K. (1988). *Multiple Representations from a Detailed Database: A Scheme for Automated Generalization*. (Ph.D. Thesis, University of Wisconsin, Madison).
- Berg, J. (1983). Aristotle's Theory of Definition. *ATTI del Convegno Internazionale di Storia della Logica 1982* (pp. 19-30). Bologna: CLUEB.
- Bernauer, J. (1994). Subsumption principles underlying medical concept systems and their formal reconstructions. *Proceedings of the Annual Symposium on Computer Applications in Medical Care* (pp. 140-144). Washington: Heinley and Belfus.
- BFO: Basic Formal Ontology (BFO). Retrieved January 16, 2008 from <http://www.ifomis.uni-Saarland.de/bfo/>.
- Bibel, W., Schneeberger, J. & Elver, E. (1990). Representation of Knowledge. *Knowledge Engineering Vol. I: Fundamentals* (pp. 1-27). New York: McGraw-Hill.
- Bittner, T. (1997). A Qualitative Coordinate Language of Location of Figures within the Ground. In S. C. Hirtle & A. U. Frank (Eds.) *Spatial Information Theory: A Theoretical Basis for GIS, International Conference COSIT '97* (pp. 223-240). Berlin: Springer.
- Bittner, T. & Smith, B. (2001). Granular Partitions, Vagueness, and Approximation. In *Proceedings of the International Conference on Formal Ontology and Information Systems - Volume 2001* (pp. 309-321). New York: ACM Press.
- Bittner, T. & Smith, B. (2003). A Theory of Granular Partitions. In M. Duckham, M. F. Goodchild & M. F. Worboys (Eds.), *Foundations of Geographic Information Science* (pp. 117-151). London: Taylor & Francis.
- Bittner, T. & Stell, J. G. (1998). A Boundary-Sensitive Approach to Qualitative Location. *Annals of Mathematics and Artificial Intelligence*, 24, 93-114.
- Bodenreider, O., Smith, B. & Burgun, A. (2004). The Ontology-epistemology Divide: A Case Study in Medical Terminology. In *Proceedings of the Third International Conference on Formal Ontology in Information Systems (FOIS 2004)* (pp. 185-195). Amsterdam: IOS Press.

- Bodenreider, O., Smith, B. & Burgun, A. (2004a). Investigating Subsumption in DL-based Terminologies: A Case Study in SNOMED CT. In U. Hahn, S. Schulz & R. Cornet (Eds.), *Proceedings of the First International Workshop on Formal Biomedical Knowledge Representation (KR-MED 2004)*. Retrieved from <http://CEUR-WS.org/Vol-102/>.
- Bonitz, H. (1967). *Über die Kategorien des Aristoteles*. Darmstadt: Wissenschaftliche Buchgesellschaft.
- Bonnet, C. (1766). *The Contemplation of Nature*. London: T. Longman.
- Borges, J. L. (1981). The Analytical Language of John Wilkins. In E. R. Monegal & A. Reid (Eds.), *Borges: A Reader* (pp. 141-143). New York: Dutton.
- Brachman, R. (1979). On the Epistemological Status of Semantic Networks. In N. Findler (Ed.), *Associative Networks: Representation and Use of Knowledge by Computers* (pp. 3-50). New York: Academic Press.
- Brentano, F. (1862). *Von der mannigfachen Bedeutung des Seienden nach Aristoteles*. Freiburg: Herder.
- Bucher, T. (1998). *Einführung in die angewandte Logik*. Berlin: De Gruyter.
- Buck, R. C. & Hull, D. L. (1966). The Logical Structure of the Linnaean Hierarchy. *Systematic Zoology*, 15 (2), 97-111.
- Bunge, W. (1966). *Theoretical Geography*. Sweden: Lund.
- Campbell, K. (1990). *Abstract Particulars*. Oxford: Blackwell.
- Campbell, K., Oliver, D. Spackman, K. & Shortliffe, E. (1998). Representing Thoughts, Words, and Things in the UMLS. *Journal of the American Medical Informatics Association*, 5 (5), 421-431.
- Caplan, A. L. (1981). Back to Class: A Note on the Ontology of Species. *Philosophy of Science*, 48, 130-140.
- Carnap, R. (1934). *The Logical Syntax of Language*. London: Routledge.
- Carroll, L. (1965). *The Works of Lewis Carroll*. London: Paul Hamlyn.
- Casati, R. & Varzi, A. C. (1994). *Holes and Other Superfacilities*. Cambridge, MA: MIT Press.
- Casati, R. & Varzi, A. C. (1995). The Structure of Spatial Location. *Philosophical Studies*, 82, 205-239.
- Casati, R. & Varzi, A. C. (1999). *Parts and Places: The Structures of Spatial Representation*. Cambridge, MA: MIT Press.
- Ceusters W. (2006). Towards a Realism-Based Metric for Quality Assurance in Ontology Matching. In B. Bennett & C. Fellbaum (Eds.), *Formal Ontology in Information Systems: Proceedings of the Fourth International Conference (FOIS 2006)* (pp. 321-332). Amsterdam: IOS Press.

- Ceusters, W., Elkin, P. & Smith, B. (2006). Referent Tracking: The Problem of Negative Findings. In *Proceedings of Medical Informatics Europe*, 38, 741-746.
- Ceusters, W. & Smith, B. (2005). Tracking Referents in Electronic Health Records. In *Proceedings of Medical Informatics Europe*, 29, 71-76.
- Ceusters, W. & Smith, B. (2006). A Realism-Based Approach to the Evolution of Biomedical Ontologies. In *Proceedings of the Annual AMIA Symposium*, Washington, DC (pp. 121-125).
- Ceusters, W. & Smith, B. (2006a). Strategies for Referent Tracking in Electronic Health Records. In *Journal of Biomedical Informatics*, 39 (3), 362-378.
- Ceusters, W. & Smith, B. (2007). Referent Tracking and its Applications. In *Proceedings of the WWW2007 Workshop i3: Identity, Identifiers, Identification*. http://ceur-ws.org/Vol-249/submission_105.pdf.
- Ceusters, W., Smith, B., Kumar, A. & Dhaen, C. (2004). Mistakes in Medical Ontologies: Where Do They Come From and How Can They Be Detected? In D. M. Pisanelli (Ed.), *Studies in Health Technology and Informatics: Ontologies in Medicine* (pp. 145-164). Amsterdam: IOS Press.
- Ceusters, W., Smith, B. & Goldberg, L. (2005). A Terminological Analysis and Ontological Analysis of the NCI Thesaurus. In *Methods of Information in Medicine*, 44, 498-507.
- Ceusters, W., Spackman, K. A. & Smith, B. Would SNOMED CT benefit from Realism-Based Ontology Evolution? In *Proceedings of the American Medical Informatics Association 2007 Annual Symposium* (pp. 105-109).
- Charlton, W. (1972). Aristotle and the Principle of Individuation. *Phronesis*, 17, 239-249.
- Chrisman, N. (1982). *Models of Spatial Analysis Based on Error in Categorical Maps*. (Ph.D. Thesis. University of Bristol, England).
- Cimino, J. J. (1998). Desiderata for Controlled Medical Vocabularies in the Twenty-first Century. *Methods of Information in Medicine*, 37 (4-5), 394-403.
- Coady, C. (1992). *Testimony: A Philosophical Study*. New York: Oxford University Press.
- Cohn, A. G. & Mark, D. M. (Eds.). (2005). *Spatial Information Theory*. In *Proceedings of COSIT 2005 (Lecture Notes in Computer Science)*. Berlin: Springer.
- Comrie, B. (1976). *Aspect*. Cambridge: Cambridge University Press.
- Cuvier, G. (1827). *The Animal Kingdom, Arranged after Its Conformity with Its Organization*. London: Whitacker.

- Cyc: OpenCyc. (n.d.). Open source version of the Cyc technology. Retrieved August 8, 2006 from <http://www.cyc.com/cyc/opencyc/overview>.
- Cyc Upper Ontology (n.d.). Retrieved August 8, 2006 from <http://www.cyc.com/cycdoc/upperont-diagram.html>.
- Davidson, D. (1980). The Logical Form of Action Sentences. *Essays on Actions and Events* (pp. 105-122). Oxford: Clarendon Press.
- Defense Mapping Agency (1980). *Production Specifications for 1: 50,000 Scale Topographic Maps of Foreign Areas*. Washington, DC: Defense Mapping Agency.
- Degen, W., Heller, B., Herre, H. & Smith, B. (2001). GOL: Towards an Axiomatized Upper-Level Ontology. In C. Welty, & B. Smith (Eds.) *Proceedings of the International Conference on Formal Ontology and Information Systems - Volume 2001* (pp. 34-46). New York: ACM Press.
- Dowty, D. R. (1991). *Word Meaning and Montague Grammar*. Dordrecht: Kluwer Academic Publishers.
- Dugas, M. & Schmidt, K. (2003). *Medizinische Informatik und Bioinformatik: Ein Kompendium für Studium und Praxis*. Berlin: Springer.
- Ebert, T. (1985). Gattungen der Prädikate und Gattungen des Seienden bei Aristoteles. Zum Verhältnis von Kat. 4 und Top. I 9. *Archiv für Geschichte der Philosophie*, 67, 113-138.
- Edelman, S. (1998). Representation is Representation of Similarities. *Behavioral and Brain Sciences*, 21, 449-498.
- Ereshefky, M. Species. *Stanford Encyclopedia of Philosophy*. Retrieved August 15, 2002 from <http://plato.stanford.edu/entries/species/>.
- Erwig, M. & Schneider, M. (1999). The Honeycomb Model of Spatio-Temporal Partitions. In M. Bohlen, C. Jensen & M. Scholl (Eds.) *International Workshop on Spatio-Temporal Database Management* (pp. 39-59). Berlin: Springer.
- Feibleman, J. K. (1974). Professor Quine and Real Classes. *Notre Dame Journal of Formal Logic*, 15, 207-224.
- Fellbaum, C. (Ed.). (1998). *WordNet: An Electronic Lexical Database*. Cambridge, MA: MIT Press.
- Foucault, M. (1970). *The Order of Things*. London: Tavistock Publications.
- Foundational Model of Anatomy (FMA). (n.d.). Retrieved January 16, 2008 from <http://sig.biostr.washington.edu/projects/fm/>.
- Fragoso, G., et al. (2004). Overview and Utilization of the NCI Thesaurus. *Comparative and Functional Genomics*, 5, 648-654.

- Frank, A., McGranaghan, M. & Volta, G. (1997). Formalization of Families of Categorical Coverages. *International Journal of Geographic Information Science*, 11, 214-231.
- Franklin, J. (2002). Stove's Discovery of the Worst Argument in the World. *Philosophy*, 77, 615-624.
- Frege, G. (1892). Begriff und Gegenstand. *Vierteljahrschrift für wissenschaftliche Philosophie*, 16, 192-205.
- Frege, G. (1980). *The Foundations of Arithmetic: A Logico-Mathematical Enquiry into the Concept of Number*. (J. L. Austin, Trans.). Oxford: Blackwell. (Original work published in 1884. *Grundlagen der Arithmetik*. Breslau: Koebner.)
- Galton, A. (1984). *The Logic of Aspect: An Axiomatic Approach*. Oxford: Clarendon Press.
- Galton, A. (1999). The Mereotopology of Discrete Space. In C. Freska & D. M. Mark (Eds.), *Spatial Information Theory: Cognitive and Computational Foundations of Geographic Science (Lecture Notes in Computer Science 1661)* (pp. 251-266). Berlin: Springer.
- Gärdenfors, P. (Ed). (2003). *Belief Revision*. Cambridge: Cambridge University Press.
- Gaus, W. (2003). *Dokumentations- und Ordnungslehre: Theorie und Praxis des Information Retrieval* (4th ed.). Berlin: Springer.
- Gelman, S. A. & Wellman, H. M. (1991). Insides and Essences: Early Understandings of the Non-obvious. *Cognition*, 38, 213-244.
- Gene Ontology Anotation Tool (GOAT). (n.d.). Retrieved January 28, 2008 from <http://goat.man.ac.uk/>.
- Gene Ontology Consortium. (2000). The Gene Ontology: Tool for the Unification of Biology. *Nature Genetics*, 25, 25-29.
- Gene Ontology Consortium. (n.d.). Retrieved January 16, 2008 from <http://www.geneontology.org>.
- Genesereth, M. & Nilsson, L. (1987). *Logical Foundations of Artificial Intelligence*. Los Altos, CA: Morgan Kaufmann.
- Ghiselin, M. (1974). A Radical Solution to the Species Problem. *Systematic Zoology*, 23, 536-544.
- Ghiselin, M. (1997). *Metaphysics and the Origin of Species*. Albany, NY: State University of New York Press.
- Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.

- Gill, K. (1993). On the Metaphysical Distinction between Processes and Events. *Canadian Journal of Philosophy*, 23, 365-384.
- Goclenius, R. (1613). *Lexicon Philosophicum*. Frankfurt.
- Goldman, A. I. (1970). *A Theory of Human Action*. Engelwood Cliffs, NJ: Prentice Hall.
- Grenon, P. (2003). Tucking RCC in Cyc's Ontological Bed. In G. Gottlob & T. Walsh (Eds.), *Proceedings of the 18th International Joint Conference in Artificial Intelligence* (pp. 894-899). San Francisco: Morgan Kaufmann.
- Grenon, P. (2003a). Spatio-temporality in Basic Formal Ontology, SNAP and SPAN, Upper Level Ontology, and Framework for Formalization, PART I. *IFOMIS Reports*. ISSN 1611-4019. Saarbrücken: Saarland University.
- Grenon, P. (2006). Temporal Qualification and Change with First-Order Binary Predicates. In B. Bennett & C. Fellbaum (Eds.), *Formal Ontology and Information Systems. Proceedings of the Fourth International Conference* (pp. 155-166). Amsterdam: IOS Press.
- Grenon, P. & Smith, B. (2004). SNAP and SPAN: Towards Dynamic Spatial Ontology. *Spatial Cognition and Computation*, 4, 69-103.
- Grenon, P., Smith, B. & Goldberg, L. (2004). Biodynamic Ontology: Applying BFO in the Biomedical Domain. In D. M. Pisanelli (Ed.), *Ontologies in Medicine: Proceedings of the Workshop on Medical Ontologies, Rome, October 2003* (pp. 20-38). Amsterdam: IOS Press.
- Griffith, P. E. (1996). Darwinism, process structuralism, and natural kinds. *Philosophy of Science*, 61, 206-227.
- Guarino, N. (2001). Review of Sowa's Knowledge Representation. *AI Magazine*, 2(3), 123-124.
- Guarino, N. & Welty, C. (2000). Ontological Analysis of Taxonomic Relationships. In A. Laender & V. Storey (Eds.) *Proceedings of ER-2000: The 19th International Conference on Conceptual Modeling* (pp. 210-224). Berlin: Springer.
- Hennig, W. (1966). *Phylogenetic Systematics*. Urbana, IL: University of Illinois Press.
- Heuer, P. (2006). *Art, Gattung, System: Logisch-systematische Analyse einiger Grundbegriffe der Biologie*. (Ph.D. Thesis. University of Leipzig).
- Hochberg, H. (1965). Universals, Particulars and Predication. *Review of Metaphysics*, 19, 87-102.
- Hochberg, H. (1966). Things and Descriptions. *American Philosophical Quarterly*, 3, 1-9.

- Hochberg, H. (1969). Moore and Russell on Particulars, Relations and Identity. In E. D. Klemke (Ed.), *Studies in the Philosophy of G. E. Moore* (pp.155-194). Chicago: Quadrangle.
- Hobbes, T. (1999). *De corpore: Elementorum philosophiae sectio prima*. Paris: Vrin.
- Hodges, W. (2001). *Logic: An Introduction to Elementary Logic*. London: Penguin Books.
- Hull, D. L. (1976). Are Species Really Individuals?. *Systematic Zoology*, 25, 174-191.
- Hull, D. L. (1997). The Ideal Species Concept and Why We Can't Get It. In M. Claridge, A. Dawah & M. Wilson (Eds.), *Species: The Units of Biodiversity* (pp. 357-380). London: Springer.
- Husserl, E. (1931). *Ideas: General Introduction to Pure Phenomenology*. (W. R. Gibson, Trans.). London: Collier-Macmillan.
- Husserl, E. (1970). *Logical Investigations Vol. 2*. London: Routledge.
- Husserl, E. (1976). *Ideen zu einer reinen Phänomenologie und Phänomenologischen Philosophie, Erstes Buch: Allgemeine Einführung in die reine Phänomenologie*. Den Haag: Nijhoff.
- Husserl, E. (1993). *Logische Untersuchungen*. Tübingen: Max Niemeyer.
- IEEE. (n.d.). Retrieved January 14, 2008 from <http://suo.ieee.org/>.
- ISO. (n.d.). Retrieved January 15, 2008 from <http://www.iso.org/iso/home.htm>.
- ISO. (1990/2000). *Vocabulary of Terminology. ISO-1087:1990 and ISO-1087:2000*. Geneva: International Organization for Standardization.
- ISO. (2005). *Guidelines for the Construction, Format and Management of Monolingual Controlled Vocabularies. ANSI/ISO Z39.19-2005* ISBN: 1-880124-65-3. Geneva: International Organization for Standardization.
- Jansen, L. (2006). Aristoteles' Kategorie des Relativen zwischen Dialektik und Ontologie. *Philosophiegeschichte und logische Analyse*, 9, 79-104.
- Jansen, L. (2007). Aristotle's Categories. *Topoi*, 26, 151-158.
- Johansson, I. (1975). *A Critique of Karl Popper's Methodology*. Gothenburg: Scandinavian University Books.
- Johansson, I. (1986). Levels of Intension and Theories of Reference. *Theoria*, 52, 1-16.
- Johansson, I. (2000). Determinables as Universals. *The Monist*, 83, 101-121.
- Johansson, I. (2004). *Ontological Investigations* (2nd ed.). Frankfurt: Ontos Verlag. (Original work published 1989. London: Routledge).

- Johansson, I. (2005). Qualities, Quantities, and the Endurant-Perdurant Distinction in Top-Level Ontologies. In G. Büchel, B. Klein & T. Roth-Berghofer (Eds.), *Proceedings of the Second International Workshop on Philosophy and Informatics*. <http://CEUR-WS.org/Vol-130/>.
- Johansson, I. (2006). Bioinformatics and Biological Reality. *Journal of Biomedical Informatics*, 39, 274-287.
- Johansson, I. (2006a). Four Kinds of 'Is_A' Relations: Genus-subsumption, Determinable-subsumption, Specification, and Specialization. In I. Johansson, B. Klein & T. Roth-Berghofer (Eds.), *Contributions to the Third International Workshop on Philosophy and Informatics Saarbrücken 2006 (IFOMIS Reports 14)* (pp. 47-62).
- Johansson, I., et al. (2005). Functional Anatomy: A Taxonomic Proposal. *Acta Biotheoretica*, 53(3), 153-166.
- Johnson, W. (1921). *Logic: Part I*. Cambridge: Cambridge University Press.
- Jullien, F. (1990). *L'Art de la Liste*. Saint-Denis: Presses Universitaires de Vincennes, Université de Paris VIII.
- Kahn, C. H. (1978). Questions and Categories. In H. Hiz (Ed.), *Questions* (pp. 227-278). Dordrecht: Reidel.
- Kalman, H. (1999). *The Structure of Knowing: Existential Trust as an Epistemological Category*. Uppsala: Swedish Science Press.
- Kant, I. (1781). *Kritik der Reinen Vernunft*. Riga: J.F. Hartknoch.
- Kant, I. (1902ff). *Werke. Akademie-Ausgabe*. Berlin: Preussische Akademie der Wissenschaften.
- Kant, I. (1950). *Critique of Pure Reason*. (N. K. Smith, Trans.). London: MacMillan.
- Kanzian, C. (Ed.). (2008). Persistence. *Ontos Verlag*, 3.
- Kenny, A. (1963). *Action, Emotion, and the Will*. London: Routledge.
- Keuth, H. (2000). *Die Philosophie Karl Poppers*. Tübingen: Mohr Siebeck. English version: (2004). *The Philosophy of Karl Popper*. Cambridge: Cambridge University Press.
- Knoepffler, Nikolaus (1999). *Forschung an menschlichen Embryonen*. Stuttgart: Hirzel.
- Kornet, D., McAllister, J. & James, W. (2005). The Composite Species Concept: A Rigorous Basis for Cladistic Practice. In T. Raydon & L. Hemerink (Eds.), *Current Themes in Theoretical Biology: A Dutch Perspective* (pp. 95-127). Dordrecht: Springer.

- Kumar, A. & Smith, B. (2005). Oncology Ontology in the NCI Thesaurus. In *Proceedings of AIME 2005: Lecture Notes in Computer Science 3581* (pp. 213-220). Berlin: Springer.
- Lenat, D. & Guha, R. (1989). *Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project*. Reading, MA: Addison-Wesley.
- Lewin, K. (1922). *Der Begriff Genese in Physik, Biologie und Entwicklungsgeschichte: Eine Untersuchung zur vergleichenden Wissenschaftslehre*. Berlin: Springer.
- Lewis, D. (1986). *On the Plurality of Worlds*. Oxford: Oxford University Press.
- Lewis, D. (1991). *Parts of Classes*. Oxford: Blackwell.
- Lindström, P. (2001). Näst sista ordet. *Filosofisk tidskrift*, 1, 3-6.
- Lorhard, J. (1613). *Theatrum philosophicum*. Basilea.
- Lowe, E. J. (2002). *A Survey of Metaphysics*. Oxford/New York: Oxford University Press.
- Lowe, E. J. (2006). *The Four-Category-Ontology: A Metaphysical Foundation for Natural Science*. Oxford: Oxford University Press.
- Macdonald, C. (1998). Tropes and Other Things. In S. Laurence & C. Macdonald (Eds.), *Contemporary Readings in the Foundations of Metaphysics* (pp. 329-350). Oxford: Basil Blackwell.
- Mark, D. M. (1978). Topological Properties of Geographic Surfaces: Applications in Computer Cartography. *Harvard Papers on Geographic Information Systems*. Cambridge, MA: Laboratory for Computer Graphics and Spatial Analysis.
- Margolis, E. & Laurence, S. (Eds.). (1999). *Concepts: Core Readings*. Cambridge MA: MIT Press.
- Mayr, E. (1969). *Principles of Systematic Zoology*. New York: McGraw-Hill.
- Mayr, E. (1996). What Is a Species, and What Is Not?. *Philosophy of Science*, 63, 262-277.
- Mayden, R. L. (1997). A Hierarchy of Species Concepts: The Denouement in the Saga of the Species Problem. *Species: The Units of Biodiversity. op. cit.*, 381-424.
- McCray, A. (2003). An Upper Level Ontology for the Biomedical Domain. *Comparative and Functional Genomics*, 4, 80-84.
- Menzel, C. (2003). Ontology Theory. In J. Euzenat, A. Gomez-Perez, N. Guarino & H. Stuckenschmidt (Eds.), *Ontologies and Semantic Interoperability*. Retrieved from <http://CEUR-WS.org/Vol-64/>.

- Meixner, U. (2005). On Negative and Disjunctive Properties. In K. Mulligan (Ed.), *Language, Truth and Ontology* (pp. 28-36). Dordrecht: Kluwer.
- Michael, J., Mejino, J. & Rosse, C. (2001). The Role of Definitions in Biomedical Concept Representation. In S. Bakken (Ed.), *Proceedings of the American Medical Informatics Association Symposium* (pp.463-467). Philadelphia: Hanley and Belfus.
- Millikan, R. (1988). *Language, Thought, and Other Biological Categories: New Foundations for Realism*. Cambridge, MA: MIT Press.
- Millikan, R. (1999). Historical Kinds and the Special Sciences. *Philosophical Studies*, 95, 45-65.
- Mislove, M., Moss, L. & Oles, F. (1990). Partial Sets. *Situation Theory and Its Applications* (pp. 117-131). Stanford, CA: Center for the Study of Language and Information.
- Mount, D. W. (2004). *Bioinformatics: Sequence and Genome Analysis* (2nd ed.). Cold Spring Harbor: Cold Spring Harbor Laboratory Press.
- Mourelatos, A. (1978). Events, Processes, and States. *Linguistics and Philosophy*, 2, 415-434.
- Mourelatos, A. (1993). Aristotle's *Kinesis/Energeia* Distinction: A Marginal Note on Kathleen Gill's Paper. *Canadian Journal of Philosophy*, 23, 385-388.
- Myrdal, G. (1968). *Objektivitetsproblemet i samhällsforskningen*. Stockholm: Rabén & Sjögren.
- Myrdal, G. (1973). *Critical Essays on Economics*. New York: Pantheon Books.
- Nagel, E. (1961). *The Structure of Science*. London: Routledge.
- Nagel, T. (1986). *The View from Nowhere*. Oxford: Oxford University Press.
- Nagel, T. (1997). *The Last Word*. Oxford: Oxford University Press.
- Myrdal, G. (1956). *Value in Social Theory: A Selection of Essays in Methodology*. London: Routledge.
- Nardi, D. & Brachman, R. J. (2003). An Introduction to Description Logics. In F. Baader, et al. (Eds.), *The Description Logics Handbook* (pp. 1-40). Cambridge: Cambridge University Press.
- National Cancer Institute. (n.d.). National Cancer Institute Thesaurus (NCIT). Retrieved November 1, 2005 from <http://ncitterms.nci.nih.gov/NCIBrowser/Dictionary>.
- National Library of Medicine. (n.d.). National Institutes of Health, Health and Human Services. Retrieved January 16, 2008 from <http://www.nlm.nih.gov/research/umls/>.

- Niles, I. & Pease, A. (2001). Towards a Standard Upper Ontology. In C. Welty & B. Smith (Eds.), *Proceedings of the 2nd International Conference on Formal Ontology in Information Systems (FOIS-2001)* (pp. 29). New York: ACM Press.
- OBO: Open Biomedical Ontologies Consortium. (n.d.). The OBO Foundry. Retrieved January 17, 2008 from <http://obofoundry.org/>.
- OBO Evaluation Criteria, Version 7 (2006). Retrieved January 14, 2006 from http://sourceforge.net/mailarchive/forum.php?forum_id=12140&max_rows=25&style=nested&viewmonth=200602
- Oehler, K. (1986). *Aristoteles: Werke in deutscher Übersetzung. Vol. 1/1: Kategorien* (2nd ed.). Berlin: Akademie Verlag.
- Olson, E. T. (1999). *The Human Animal: Personal Identity without Psychology*. Oxford: Oxford University Press.
- Ontologyportal: Suggested Upper Merged Ontology (SUMO) Ontology Portal. (n.d.). Retrieved January 2, 2008 from <http://www.ontologyportal.org/>.
- Ouzounis, C. A. (2003). Classification Schemes for Protein Structure and Function. *Nature Reviews Genetics*, 4, 508-519.
- Peirce, C. S. (1998). *Collected Papers*. Bristol: Thoemmes.
- Perl Foundation. (n.d.). Retrieved January 15, 2008 from www.perl.org.
- Pisanelli, D. M. (Ed.). (2004). *Ontologies in Medicine: Proceedings of the Workshop on Medical Ontologies, Rome, October 2003*. Amsterdam: IOS Press.
- Plato. (1961). Phaedrus. In E. Hamilton & H. Cairns (Eds.), *Plato: The Collected Dialogues* (pp. 475-526). Princeton: Princeton University Press.
- Plato. (1961a). Theaetetus. In E. Hamilton & H. Cairns (Eds.), *Plato: The Collected Dialogues* (pp. 845-919). Princeton: Princeton University Press.
- Popper, K. (1945). *The Open Society and Its Enemies*. London: Routledge and Kegan Paul.
- Popper, K. (1959). *The Logic of Scientific Discovery*. London: Hutchinson.
- Popper, K. (1963). *Conjectures and Refutations: The Growth of Scientific Knowledge*. London: Routledge.
- Popper, K. (1972). *Objective Knowledge*. Oxford: Clarendon Press.
- Pseudo-Dionysius Areopagita. (1857-1866). *De divinis nominibus, Patrologia Graeca*, 585-984. J. P. Migne (Ed.). Paris.
- Putnam, H. (1975). The Meaning of 'Meaning'. *Mind, Language and Reality* (pp. 215-271). Cambridge: Cambridge University Press.
- Rea, M. C. (1995). The Problem of Material Constitution. *The Philosophical Review*, 104(4), 525-552.

- Rector, A., Nolan, W. & Kay, S. (1991). Foundations for an Electronic Medical Record. *Methods of Information in Medicine*, 30, 179-186.
- Reydon, T. A. (2005). *Species as Units of Generalization in Biological Science: A Philosophical Analysis*. (Ph.D. Thesis. University of Leiden).
- Ridder, L. (2002). *Mereologie: Ein Beitrag zur Ontologie und Erkenntnistheorie*. Frankfurt: Klostermann.
- Rödl, S. (2005). *Kategorien des Zeitlichen: Eine Untersuchung der Formen des endlichen Verstandes*. Frankfurt: Suhrkamp Verlag.
- Rosse, C. & Mejino, J. (2003). A Reference Ontology for Bioinformatics: The Foundational Model of Anatomy. *Journal of Biomedical Informatics*, 36, 478-500.
- Rosse, C. & Mejino, J. (2007). The Foundational Model of Anatomy Ontology. In A. Burger, D. Davidson & R. Baldock (Eds.), *Anatomy Ontologies for Bioinformatics: Principles and Practice*. New York: Springer.
- Russell, B. (1940). *An Inquiry into Meaning and Truth*. London: George Allen and Unwin.
- Russell, B. (1948). *Human Knowledge: Its Scope and Limits*. London: George Allen and Unwin.
- Russell, B. (1959). *My Philosophical Development*. London: George Allen and Unwin.
- Russell, S. & Norvig, P. (1995). *Artificial Intelligence: A Modern Approach*. Englewood Cliffs, NJ: Prentice Hall.
- Ryle, G. (1949). *The Concept of Mind*. Chicago: University of Chicago Press.
- Schark, M. (2005). *Lebewesen versus Dinge: Eine metaphysische Studie*. Berlin: De Gruyter.
- Schuhmann, K. & Smith, B. (1987). Questions: An Essay in Daubertian Phenomenology. *Philosophy and Phenomenological Research*, 47, 353-384.
- Searle, J. (1959). Determinables and the Notion of Resemblance. *Proceedings of the Aristotelian Society*, 33, 141-158.
- Searle, J. (1983). *Intentionality*. Cambridge: Cambridge University Press.
- Sellars, W. F. (1963). Philosophy and the Scientific Image of Man. *Science, Perception and Reality* (pp. 1-40). London: Routledge.
- Simons, P. M. (1987). *Parts: A Study in Ontology*. Oxford: Clarendon Press.
- Simons, P. M. (1992). Categories and Ways of Being. In *Philosophy and Logic in Central Europe from Bolzano to Tarsk: Selected Essays* (pp. 377-394). Dordrecht: Reidel.
- Simons, P. M. (2005). Against Set Theory. In M. E. Reicher & J. C. Marek (Eds.), *Experience and Analysis* (pp. 143-152). Vienna: HPT&ÖBV.

- Smith, B. (1991). Relevance, Relatedness and Restricted Set Theory. In G. Schurz & G. Dorn (Eds.), *Advances in Scientific Philosophy: Essays in Honour of Paul Weingartner* (pp. 45-56). Amsterdam: Rodopi.
- Smith, B. (1999). Truthmaker Realism. *Australasian Journal of Philosophy*, 77(3) 274-291.
- Smith, B. (2001). Fiat Objects. *Topoi*, 20(2), 131-148.
- Smith, B. (2001a). True Grid. In D. R. Montello (Ed.), *Spatial Information Theory: Foundations of Geographic Information Science* (pp. 14-27). Berlin: Springer.
- Smith, B. (2003). Ontology. In L. Floridi (Ed.), *Blackwell Guide to the Philosophy of Computing and Information* (pp. 155-166). Oxford: Wiley.
- Smith, B. (2003a). Aristotels 2002. In T. Buchheim, H. Flashar & R. King (Eds.), *Kann man heute noch etwas anfangen mit Aristoteles?* (pp. 3-38). Darmstadt: Meiner Felix Verlag.
- Smith, B. (2004). Beyond Concepts: Ontology as Reality Representation. In A. Varzi & L. Vieu (Eds.), *Formal Ontology and Information Systems*, (pp. 73-84). Amsterdam: IOS Press.
- Smith, B. (2005). The Logic of Biological Classification and the Foundations of Biomedical Ontology. In P. Hájek, et al. (Eds.), *Logic, Methodology and Philosophy of Science* (pp. 505-520). London: King's College Publications.
- Smith, B. (2005a). Against Fantology. In M. Reicher & J. Marek (Eds.), *Experience and Analysis* (pp. 153-170). Vienna: HPT&ÖPV.
- Smith, B. (2006) Against Idiosyncrasy in Ontology Development. In B. Bennett & C. Fellbaum (Eds.), *Formal Ontology in Information Systems* (pp. 15-26). Amsterdam: IOS Press.
- Smith, B., et al. (2005). Relations in Biomedical Ontologies. *Genome Biology*, 6, R46.
- Smith B., et al. (2007). The OBO Foundry: Coordinated Evolution of Ontologies to Support Biomedical Data Integration. *Nature Biotechnology*, 25, 1251-1255.
- Smith, B. & Brogaard, B. (2002). Quantum Mereotopology. *Annals of Mathematics and Artificial Intelligence*, 35 (1-2), 153-175.
- Smith, B. & Brogaard, B. (2003). A Unified Theory of Truth and Reference. *Logique et Analyse*, 43, 49-93.
- Smith, B. & Ceusters, W. (2007). Ontology as the Core Discipline of Biomedical Informatics: Legacies of the Past and Recommendations for the Future Direction of Research. In G. Dodig-Crnkovic & S. Stuart (Eds.), *Computation, Information, Cognition: The Nexus and The Liminal* (pp. 104-122). Newcastle: Cambridge Scholars Press.

- Smith, B., Ceusters, W. & Siebert, D. (2004). Was die philosophische Ontologie zur biomedizinischen Informatik beitragen kann. *Information: Wissenschaft und Praxis*, 55, 143-146.
- Smith, B., Ceusters, W. & Temmerman, R. (2005). Wüsteria. In *Medical Informatics Europe (MIE 2005)*, Geneva, *Studies in Health Technology and Informatics* (pp. 647-652). Amsterdam: IOS Press.
- Smith B. & Grenon, P. (2004). The Cornucopia of Formal-Ontological Relations. *Dialectica*, (58) 2004: 279-296
- Smith, B., Köhler, J. & Kumar, A. (2004). On the Application of Formal Principles to Life Science Data: A Case Study in the Gene Ontology. In E. Rahm (Ed.), *Proceedings of Data Integration in the Life Sciences (DILS 2004)* (pp. 79-94). Berlin: Springer.
- Smith, B., Kusnierczyk, W., Schober, D. & Ceusters, W. (2006). Towards a Reference Terminology for Ontology Research and Development in the Biomedical Domain. In O. Bodenreider (Ed.), *Proceedings of the Second International Workshop on Formal Biomedical Knowledge Representation (KR-MED 2006)* (pp. 57-65). Retrieved from <http://www.CEUR-WS.org/Vol-222/>.
- Smith, B. & Mark, D. M. (1999). Ontology with Human Subjects Testing: An Empirical Investigation of Geographic Categories. *American Journal of Economics and Sociology*, 58 (2), 245-272.
- Smith B., Mejino, J., Schulz, S., Kumar, A. & Rosse, C. (2005). Anatomical Information Science: Spatial Information Theory. *Proceedings of COSIT 2005* (pp. 149-164). Berlin: Springer.
- Smith, B. & Mulligan, K. (1983). Framework for Formal Ontology. *Topoi*, 2, 73-85.
- Smith, B., Munn, K. & Papakin, I. (2004). Bodily Systems and the Spatial-Functional Structure of the Human Body. In D. Pisanelli (Ed.), *Ontologies in Medicine: Proceedings of the Workshop on Medical Ontologies* (pp. 39-63). Amsterdam: IOS Press.
- Smith, B. & Rosse, C. (2004). The Role of Foundational Relations in the Alignment of Biomedical Ontologies. In M. Fieschi, et al. (Eds.), *Medinfo 2004* (pp. 444-448). Amsterdam: IOS Press.
- Smith, B. & Smith, D. W. (1995). Introduction. In B. Smith & D. W. Smith (Eds.), *The Cambridge Companion to Husserl*. Cambridge: Cambridge University Press.
- Smith, B. & Varzi, A. (1999). The Niche. *Noûs*, 33, 198-222.
- Smith, B. & Varzi, A. (2002). Surrounding Space: The Ontology of Organism-Environment Relations. *Theorie in den Biowissenschaften*, 121, 139-162.

- Smith, B., Williams, J. & Schulze-Kremer, S. (2003). The Ontology of the Gene Ontology. In *Proceedings of the Annual Symposium of the American Medical Informatics Association* (pp. 609-613). Amsterdam: IOS Press.
- SNOMED (The Systematized Nomenclature of Medicine). (n.d.). Retrieved August 6, 2006 from <http://www.snomed.org/snomedct/>.
- Sowa, J. (n.d.). Semantic Networks. Retrieved June 2, 2006 from <http://www.jfsowa.com/pubs/semnet.htm>.
- Sowa, J. F. (2000). *Knowledge Representation: Logical, Philosophical and Computational Foundations*. Pacific Grove, CA: Brooks Cole.
- Sowa, J. F. (2001). Top-Level Categories. Retrieved August 8, 2006 from <http://users.bestweb.net/~sowa/ontology/toplevel.htm>.
- Stekeler-Weithofer, P. (2000). Kritik der Erkenntnistheorie: Zur Logik von Gegenstandsbezug und Wahrheit bei Hegel (und Wittgenstein). In C. Iber (Ed.), *Hegels Seinslogik: Interpretationen und Perspektiven* (pp. 59-79). Berlin: Akademie Verlag.
- Stell, J. G. (1999). Granulation for Graphs. In C. Freksa & D. M. Mark (Eds.), *Spatial Information Theory: Cognitive and Computational Foundations of Geographic Information Science (Lecture Notes in Computer Science, 1661)* (pp. 417-432). Berlin: Springer.
- Stell, J. G. (2000). The Representation of Discrete Multi-Resolution Spatial Knowledge. In *Proceedings of Seventh International Conference on Principles of Knowledge Representation and Reasoning (KR2000)* (pp. 38-49). San Francisco: Morgan Kaufmann Publishers.
- Stell, J. G. & Worboys, M. F. (1999). Generalizing Graphs Using Amalgamation and Selection. In R. Gutting, D. Papadial & F. Lochovsky (Eds.), *Advances in Spatial Databases 1999, Proceedings* (pp. 19-32). Berlin: Springer.
- Stoecker, R. (1999). *Der Hirntod: ein medizinethisches Problem und eine moralphilosophische Transformation*. Freiburg: Karl Alber Verlag.
- Stout, R. 1997 Processes. *Philosophy*, 72, 19-27.
- Strawson, P. F. (1959). *Individuals: An Essay in Descriptive Metaphysics*. London: Routledge.
- Suárez, F. (1856). *Disputationes Metaphysicae, Opera Omnia*. Paris: Vivès.
- Taxonomicon. (n.d.). Retrieved January 15, 2008 from <http://taxonomicon.taxonomy.nl/Taxon tree.aspx>.
- Temmerman R. (2000). *Towards New Ways of Terminology Description*. Amsterdam: Benjamins.
- Thompson, M. (1995). The Representation of Life. In R. Hursthouse, et al. (Eds.), *Virtues and Reasons* (pp. 247-297). Oxford: Clarendon Press.

- Thompson, M. (2008). Naive Action Theory. *Life and Action*. Harvard University Press.
- Trettin, K. (2000). Ontologie individueller Qualitäten. In R. Hüntelmann & E. Tegtmeier (Eds.), *Ontologie* (pp. 145-164). St. Augustin: Academica.
- Vaihinger, H. (1924). *The Philosophy of As-If*. London: Routledge.
- Vendler, Z. (1972). *Linguistics in Philosophy*. Ithaca, NY: Cornell University Press.
- Vizenor, L. (2004). Actions in Health Care Organizations: An Ontological Analysis. In M. Fieschi, et al. (Eds.), *MedInfo 2004* (pp. 1403-1407). Amsterdam: IOS Press.
- Volta, G. & Egenhofer, M. (1993). Interaction with GIS Attribute Data Based on Categorical Coverages. In A. Frank & I. Campari (Eds.), *Spatial Information Theory: A Theoretical Basis for GIS* (pp. 215-233). London: Springer.
- Von Wachter, D. (2000). *Dinge und Eigenschaften: Versuch zur Ontologie*. Dettelbach: Röhl.
- Williams, D. C. (1953). The Elements of Being. *Review of Metaphysics*, 7, 171-192.
- Wilson, J. (1999). *Biological Individuality: The Identity and Persistence of Living Entities*. Cambridge: Cambridge University Press.
- Wolff, C. (1736). *Philosophia Prima sive Ontologia*. Leipzig.
- WordNet (n.d.) Retrieved January 2, 2006 from <http://www.cogsci.princeton.edu/cgi-bin/webwn2.0?stage=1&word=cat>.
- Wüster, E. (1979). *Einführung in die Allgemeine Terminologielehre und Terminologische Lexikographie*. Vienna: Springer.
- Wüster, E. (2003). The Wording of the World Presented Graphically and Terminologically. *Terminology*, 9 (2), 269-297.
- Zaiss, A., et al. (2005). Medizinische Dokumentation, Terminologie und Linguistik. In T. M. Lehmann (Ed.), *Handbuch der Medizinischen Informatik* (pp. 89-143). Vienna: Hanser.

Index

- a posteriori** (knowledge, judgment, belief, concept) 50, 51
- a priori** (knowledge, judgment, belief, concept) 50-52, 118, 201
- abstraction / abstract / to abstract** 9, 11, 14-16, 40-42, 66, 71, 84, 87, 90, 94, 97, 100, 101, 119, 190, 191, 194-196, 205, 206, 211, 242, 247, 249, 255, 257, 263, 279, 297-299, 306
- mathematical abstraction 41
 - ontological abstraction 41-43
 - physical 41-42
 - transnatural 41-42
- accident** 35, 181, 182, 186-188
- adequatism (realist adequatism)** 80
- Aertsen, J. A.** 42
- algebraism / algebraic** 75-78
- Allen, R.** 276
- ambiguous / ambiguity** 65, 89, 97, 104, 117, 136, 160, 177, 205, 219, 222
- ancestry (as criterion for biological classification)** 214-317
- Angelelli, I.** 186
- Anscombe, E.** 263
- application ontology (see ontology, application)**
- Aquinas, T.** 47, 180
- Aristotle / Aristotelian** 17, 22, 23, 26, 30, 31, 35, 36, 39, 40, 42, 46-48, 111, 132, 159, 162, 164, 167, 173, 175, 178-181, 183, 185, 187, 193-199, 207, 217, 219, 225, 243, 244, 257-259, 264, 268, 280
- Armstrong, D. M.** 23, 177, 186
- artifact** 97, 99, 143, 170, 171, 184, 203, 212, 305
- artificial intelligence / AI** 61
- Asplund, J.** 252
- attribute** 22, 34, 41-43, 51, 55, 114, 120, 168, 177, 224, 227
- Baader, F.** 236
- Baier, A. C.** 269
- Barker, R. G.** 25
- being (see existence or entity)**
- belief** 8, 9, 78, 79, 81, 92-96, 104, 105, 307, 308
- Berg, J.** 244
- Bernauer, J.** 248
- BFO (Basic Formal Ontology)** 17, 174, 175, 196, 204, 281
- biasism** 285-287, 289, 290, 291, 296, 299
- Bibel, W.** 63
- biological classification (see classification, biological)**
- biological process (only in reference to the GO upper-level category)** 36, 304-305
- biology / biological** 11, 13, 15, 17, 19, 21-36, 54, 77, 80, 88, 94, 105, 126, 128, 131, 133, 142, 144, 157, 160, 169-171, 198, 200-207, 211-217, 219, 220, 221, 223, 225-227, 229-235, 384, 285, 290, 291-296, 303
- biomedical informatics (see informatics, biomedical)**
- biomedicine (see domain, biomedical)**
- Bittner, T.** 16, 125, 127, 134

- Bodenreider, O. 91, 103
 bona fide (see object, bona fide)
 Bonitz, H. 179
 Bonnet, C. 202, 216
 Borges, J. L. 159, 163, 168, 171
 Brachman, R. 114
 Brentano, F. 180, 306
 Bucher, T. 189
 Buck, R. C. 215
- Campbell, K. 98-99, 101, 186
 Caplan, A. L. 208
 Carnap, R. 236
 Carroll, L. 181
 Casati, R. 134, 225
 category (see also class, kind, species, type, universal) 10, 13, 14, 17, 18, 23, 26, 28, 41, 67-69, 95, 164, 178-186, 192, 195
 cause / causation / causal 8, 13, 32, 40, 49, 51, 52, 54, 90, 103, 111
 causal relation (see relation, causal)
 cellular component (only in reference to the GO upper-level category) 36, 304
 Ceusters, W. 20, 96, 106-107, 116, 119-120, 123, 144, 162, 183, 192, 219, 223
 Charlton, W. 47
 Cimino, J. 88-90, 104, 107, 117
 characteristic 31, 49, 86, 87, 100, 102, 162, 167-169, 178, 189, 197, 207, 214, 224, 233, 256
 class (see also category, kind, species, type, universal) 22, 23, 30, 31, 34, 70, 71, 94, 96, 97, 126, 152, 156, 163, 169, 170, 173, 175, 177, 188, 192, 198, 200, 202, 205, 206, 214, 215, 229, 322, 236-257, 306
 classification / classify 16-18, 22, 23, 28, 30, 96, 97, 122, 125, 126, 132, 139, 144, 148, 159-170, 173-175, 177-179, 184, 193, 194, 197, 200, 202-205, 211-215, 217, 235, 252, 253
 – biological classification 17, 148, 160
 Coady, C. 291
 cognition / cognitive act 11, 12, 48, 49, 52, 54, 72, 78, 79, 87, 88, 91, 98, 125-127, 143, 148, 157, 170
 completeness / complete / incompleteness 77, 137, 143, 144, 146-149, 152-156, 170, 171, 188, 197, 198, 202, 206, 214, 215, 243, 280
 – of occurrents 12, 22, 259-264, 266, 267, 272, 273, 276, 280-283
 Comrie, B. 262
 concept / conceptual 11, 15, 18, 19, 29, 40, 48, 49, 51-55, 58, 70-117, 123, 211, 215-217, 235, 236, 238, 240, 243, 244, 246-250, 279, 281, 285-287, 294, 295, 297, 298, 300-308
 – concept logic 113-114
 – concept(ual) orientation 11, 15, 83, 84, 87-93, 96, 98, 101, 102, 107, 117
 – conceptual system 70, 302
 – conceptualism 70-78, 92, 240
 – pragmatist conceptualism 70, 73, 75, 78
 conceptualization (see representation)
 consistency / consistent / inconsistency / inconsistent 14,

- 28-31, 73, 77, 114, 136, 137, 177, 222, 223, 299
- construct (noun) / construction (of reality) 71, 76, 78
- content 8, 21, 25, 28, 70, 76, 98, 293, 297-299, 307
- formal 46, 49, 50, 51, 56
 - material 49, 50, 51, 56
- continuant 18, 32, 69, 113, 164, 173, 175, 180, 183-186, 188, 194-196, 219, 221, 225-231, 256, 257, 259, 267, 268, 276, 281
- controlled vocabulary (see terminology)
- CUI (Concept Unique Identifier) 99-100
- cumulativeness / cumulative 7, 86, 107, 146-153, 157
- Cuvier, G. 215-216
- Cyc 22, 62, 70, 73, 74, 78, 167, 174, 175, 177, 196
- Darwin, C.** 22, 207, 211, 242
- data 7, 9, 10, 21, 26, 28, 29, 35, 60, 61, 63, 64, 83, 92, 95, 105-107, 114, 120-124, 168-171, 205, 221, 282, 290
- database 8, 9, 10, 15, 21, 26, 36, 61, 81, 149, 159, 169, 170, 221
- Davidson, D. 111, 186
- definition
- Aristotelian 243-244
 - formal 115, 228, 225
 - intensional 102
 - rigorous 21, 222, 223, 225, 233
- demonstrative expression (see expression, demonstrative)
- dependent entity (see entity, dependent)
- Descartes, R. 16
- description logic (DL) (see logic, description)
- determinable-subsumption (see relation, is_a)
- disjunctive / disjunctiveness (of classification system) 163, 167
- disposition 32, 33, 177, 184, 241
- domain 11-19, 21-25, 28, 31, 34, 39, 44, 46, 54, 60-62, 73, 77, 79, 85, 88, 90, 92, 94, 95, 101, 102, 105, 107, 118, 120, 123, 125, 138, 140-146, 149, 153, 158, 160, 169, 170, 174, 185, 197, 198, 201, 202, 204, 211, 212, 219, 220, 222, 224, 227, 234, 243, 289, 291, 296, 303, 309
- biomedical domain 22, 26, 27, 84, 104, 197
 - biological domain 19, 24, 34, 203, 204, 220, 221, 223
 - domain-neutral / domain-specific 14, 67-69, 73, 224
- Dowty, D. R. 262
- Dugas, M. 167, 170
- Ebert, T.** 179
- Edelman, S. 81
- EHR (electronic health record) 95, 105, 106, 113, 124, 119
- EHR system 9, 105-107, 109, 124
- empirical 10, 31, 32, 54, 56, 77, 87, 91, 96, 104, 160, 191, 198, 200, 201, 203, 204, 212, 215, 216, 288, 295-298
- inquiry 79-81
 - object 49-50
- empty / emptiness (with reference to granular partition theory) 129,

- 136, 137, 142, 145, 146, 148, 150-157
- endurant (see continuant)
- entity 14, 16, 25, 33, 39-45, 64, 66, 69, 75, 81, 93, 97, 106, 107, 112, 113, 122, 123, 127, 133, 148, 152, 162, 170, 181, 183, 185, 191, 203, 221, 226, 228, 229, 246, 279, 299
- dependent, independent 246
- environment 15, 24, 25, 27, 35, 149
- epistemology / epistemological / epistemic 19, 22, 42, 43, 54, 78, 79, 84, 242, 285, 291-293, 295, 302
- Ereshefky, M. 207
- evolution, theory of 31-34, 36, 197, 207, 209-211, 214-217
- exemplification, exemplify (see relation, exemplification)
- exhaustiveness / exhaustive 17, 144-147, 149, 153, 160, 163, 167-169, 175, 184, 208, 242
- existence / exist 129, 154, 155, 165, 169, 171, 177, 181-185, 187, 191, 193, 196, 202-206, 213, 214, 221, 225, 228-230, 240, 243, 255, 256, 275, 280, 282, 283, 285, 293, 204, 297, 299, 307
- expert system (see system, expert)
- explicit / explicitness (of classification system) 21, 24, 83, 161, 168
- expression 21, 75, 78, 89, 94-97, 100, 106, 111, 162, 170, 177, 179, 185, 186, 189, 192, 193, 196, 219-228, 233-237, 247, 270-271, 277, 278, 297, 299
- demonstrative 45, 47, 48, 55, 56
- fact 21, 34, 56, 59, 73, 74, 78, 95, 96, 110, 139, 209-212, 216, 217, 286, 288, 290-295, 300-302, 306
- fact vs. rule 59-65, 69
- fallibilism 13, 79, 292, 293, 296
- realist fallibilism 8
- falsification, falsifiable 294, 298, 299
- Feibleman, J. K. 192, 238
- Fellbaum, C. 223
- fiat object (see object, fiat)
- First Philosophy 39
- FMA (Foundational Model of Anatomy) 10, 30, 31, 103, 115, 131, 133
- form / formal
- definition (see definition, formal)
- discipline 44-46, 50
- formal content (see content, formal)
- formal ontology (see ontology, formal)
- form of life 207-210, 213
- Foucault, M. 24, 171
- four-dimensional / four-dimensionalism 281-283
- Fragoso, G. 162
- Frank, A. 142
- Franklin, J. 93
- Frege 109, 114, 162, 168, 177, 219
- Fregean logic (see logic, Fregean)
- full / fullness (with reference to granular partition theory) 148-156
- function 15, 26, 27, 29-33, 36, 136, 137, 142, 144, 145, 147, 167, 197, 221, 224, 251, 264, 296, 304

GALEN 115, 299

- Galton, A. 262-263
- Gärdenfors, P. 96
- Gaus, W. 170
- Gelman, S. A. 88
- Gene Ontology (GO) 10, 22, 27, 29, 35, 36, 83, 221, 223, 250, 251, 254, 264, 296, 304, 305
- general / generality / generalization 9, 10, 14, 15, 17, 27, 33, 34, 41-47, 51-54, 60, 65-68, 73, 76, 84-86, 91, 92, 96, 105, 106, 110-113, 126, 134, 168, 173, 185-187, 192, 193, 197, 202-204, 216, 235, 238, 240, 241, 256-258, 264, 271, 272, 278, 289, 292, 308
- general term 27, 45, 76, 84, 92, 96, 106, 111, 113, 192, 219-224, 240
- genetic / genetics / genetic research / gene 24, 26, 27, 31-36, 88, 94, 102, 120, 162, 167, 169, 209, 210, 212, 223, 224
- Genesereth, M. 71-72
- genus 18, 30, 31, 95, 149, 164, 198, 202-205, 215, 217, 235, 242-244, 247, 249, 250
- genus-subsumption (see relation, *is_a*)
- geometry / geometrical 41, 44, 47, 201, 255, 294
- Ghiselin, M. 34, 207-208
- Gibson, J. J. 25
- Gill, K. 263
- Goclenius, R. 40-43
- Goldman, A. I. 276
- granular / granularity / level of granularity / granular level 12, 24-26, 33, 138, 144, 158, 303
- granular partition 13, 16, 125-127, 133, 134, 141, 143, 155-158
- granular partition theory 125, 134, 192
- Grenon, P. 15, 20, 32, 57, 65, 76, 174, 185, 204, 226, 281
- Griffith, P. E. 210
- Guarino, N. 195
- Guha, R. 62
- Health Level 7 Reference Information Model (HL7 RIM)** 205
- Hennig, W. 15, 17-18, 39, 197, 210, 214-215, 255
- Heuer, P. 17, 197
- hierarchy / hierarchical structure / taxonomy (see taxonomy)
- Hochberg, H. 186
- Hobbes, T. 203
- Hodges, W. 225
- Hull, D. L. 215
- Husserl, E. 43, 44, 47, 53-56, 67, 68, 307, 308
- identity** 18, 32-34, 69, 123, 140, 156, 186, 190, 203, 224
- synchronic / diachronic 34
- IEEE 73, 77
- incompleteness (of occurrents, see completeness)
- inconsistent / inconsistency (see consistent / consistency)
- individual 9, 11, 26, 33, 34, 36, 77, 85, 86, 100, 102, 109-114, 116, 126, 142, 146, 175, 177, 180, 182, 184, 186-188, 192, 198, 201, 203, 208-210, 213, 215-217, 227, 228, 237, 241, 246, 251, 253, 255, 268, 290
- individuation, principle of 47-48

- inference / infer / rule of inference
59-63, 72, 83, 95, 106, 140, 144,
170, 178, 216, 234, 305, 308
- inference engine 55, 62
- informatics / information science
28, 83, 84, 98, 173, 196, 253
- bioinformatics 15, 18, 19, 22, 26,
27, 34, 285, 287, 296
 - biomedical 28, 84, 285
- information 7-27, 50, 51, 53, 63,
64, 69, 71, 83, 87, 95, 97, 102,
106, 107, 109, 113, 115, 116, 119,
120-123, 162, 164, 174, 199, 205,
221, 222, 229, 232, 235, 236,
245, 247, 254, 291, 300, 302,
303, 305, 309
- information-gathering (see
knowledge acquisition)
- inherence, inhere (see relation,
inherence)
- inheritance, single / multiple 18,
165, 175, 236, 249-252
- instance / instantiation / instantiate
17, 18, 30-34, 39, 41, 46, 69, 77,
86, 93, 104-107, 111-124, 128,
185, 192, 193, 200-202, 205, 206,
211, 213, 215, 217, 220-235, 238-
242, 246, 252, 253, 255, 256,
258, 260, 261-263, 266, 267-270,
272-282
- instance vector 120-121
- instance_of relation (see relation,
instantiation)
- interoperable / interoperability 10,
15-17, 74, 78, 171, 221, 234
- intuition 48-52
- is_a (see relation, is_a)
- ISO (International Organization for
Standardization) 71, 73, 84-87,
102, 113, 114
- IUI (instance-unique identifier),
instance identifier 107, 121-123,
- Jansen, L.** 16-17, 159, 173, 177-
178, 180
- Johansson, I.** 18-19, 32-33, 184,
235-236, 241, 245, 285, 292
- Johnson, W.** 183, 241
- judgment 44-52, 56, 294, 295
- specific 45, 46
 - particular 45, 46
- Jullien, F.** 171
- Kahn, C. H.** 167, 179
- Kalman, H.** 291
- Kant, I.** 21, 22, 45-54, 113, 180,
200-205, 213, 257
- kind (see also class, category,
species, type, universal) 11, 15,
16, 22, 23, 26, 31, 33, 34, 42, 46,
66, 67, 72, 75, 84, 87, 88, 90, 91,
101-103, 105, 106, 111, 114, 134,
162-164, 168, 169, 173, 193, 197,
198, 200, 204, 205, 211, 215,
216, 229, 235, 237, 240-242, 244,
247, 257
- Kenny, A.** 264
- Keuth, H.** 294
- Knoepffler, N.** 23
- knowledge 7-9, 11, 26, 44, 51, 54,
56, 60, 66, 79, 83, 87, 89, 97,
104, 105, 107, 113, 114, 116, 118,
123, 149, 154, 155, 198, 210,
204, 221, 285, 291, 296, 302
- Acquisition / information
gathering 69, 106, 204
 - knowledge base 62, 64, 66, 71,
73, 74, 81

- knowledge engineer / knowledge engineering 11, 16, 63-66, 69, 71, 73, 80, 81, 291
- knowledge management 15, 16, 34, 61, 63, 205, 216
- knowledge representation 14, 15, 57-59, 61, 63, 64, 69, 70, 72-81, 83, 92, 97, 105, 114, 170, 173, 179, 196, 199, 221, 223, 235
- Kornet, D. 210
- Kumar, A. 28, 36, 165
- language**
 - artificial, first-order, formal, logical, symbolic 71, 109, 110, 117
 - natural 10, 63, 75, 76, 107, 235, 236
 - programming 18, 169, 235, 271
- Leibniz, G. W. 25, 99, 180
- Lenat, D. 62
- Lewin, K. 34
- Lewis, D. 23, 126, 130, 183, 184
- life form (see form of life)
- Lindström, P. 308
- linguisticism / linguistic 63, 75, 78, 84, 86, 90, 189, 193, 195, 221, 235, 247, 270, 271, 300-302, 306
- Linnaeus, C. 22, 23, 211
- located_in relation / location relation (see relation, location)
- logic 23, 29, 30, 44, 61-63, 73, 74, 76, 94, 99, 109, 111, 113-115, 127, 169, 177, 201, 216, 223, 232, 279, 286, 288
 - concept logic 113, 114, 123
 - description logic (DL) 65, 114, 235, 236
 - first-order logic (FOL) 65, 109, 110, 111, 113, 127, 179, 186
 - formal logic 127
 - realist logic 109, 111, 113, 116, 123
- Lowe, E. J. 184-186
- Macdonald, C.** 186
- mapping / map 12, 49, 73, 89, 107, 117, 120, 123-126, 133-135, 137, 138, 141, 142, 145, 152, 171, 281-283
- Mark, D. M. 132
- Margolis, E. 98
- mass term / mass noun 206
- material content (see content, material)
- material ontology (see ontology, material)
- mathematical abstraction (see abstraction, mathematical)
- matter / material 23, 24, 27, 33, 40-47, 68, 86, 133, 177, 203, 225, 235, 236, 293, 300-302
 - material ontology (see ontology, material)
 - matter vs. form (see form)
- Mayr, E. 207, 214
- Mayden, R. L. 207
- McAllister, J. 21
- McCray, A. 29, 285
- meaning (of concept) 11, 23, 71, 72, 55, 86, 88-90, 97, 113, 201, 285, 306
- Meixner, U. 219
- Mejino, J. 30
- mention / use distinction (see use/mention distinction)
- Menzel, C. 70

- mereology / mereological 126, 130-132, 137, 140-145, 149-156, 188, 191, 192, 228
 mereological structure 138-140, 143-145, 151, 153, 157, 158
 metaphysics / metaphysical 12, 21-23, 40, 66, 69, 80, 178, 180, 181, 259, 280, 292
 methodology / methodological 15, 21, 26, 35, 36, 57, 58, 63, 70, 72, 75, 79-81, 123, 171, 222, 223, 234, 292, 296
 Michael, J. 30
 Millikan, R. 23, 33, 210
 model 10, 16, 30, 70-72, 76, 81, 93, 97, 103, 114, 125, 132, 305
 modular / modularity 24, 30, 33
 molecular function 36, 251, 296, 304
 Mourelatos, A. 264
 Myrdal, K. G. 285-287, 290, 293

 Nagel, E. 286
 Nagel, T. 308
 name / proper name 45, 49, 106, 107, 111, 117, 118, 133, 134, 206, 210, 222, 258
 Nardi, D. 114
 natural language (see language, natural)
 NCI (National Cancer Institute), NCI Thesaurus / NCIT 16, 159, 162-165, 167, 168, 170, 172
 Niles, I. 192
 Nilsson, L. 71-72
 nominalism 240
 Norvig, P. 69

object 11, 22, 25-27, 39, 41, 43-46, 48-55, 63, 64, 68, 71, 76, 77, 85-88, 90, 91, 94-96, 100, 102, 110, 113, 115, 125-128, 131-149, 151-153, 155, 158, 160, 161, 163, 177, 195, 210, 210, 212, 213, 221, 235, 240, 248, 249, 258, 281, 282, 306, 307
 – bona fide / fiat 126, 245
 – object as such 44, 56
 – of experience 44, 49, 50, 54-56
 – of knowledge 11, 22, 24, 57
 OBO (Open Biomedical Ontologies) 17, 29, 35, 37, 104, 171, 174, 196
 – OBO Evaluation Criteria 170
 – OBO Foundry 29, 30, 35, 37, 104, 171, 196
 observation (clinical, empirical etc. ...) 51, 95, 106, 107, 253
 occurrent 18, 32, 46, 111, 113, 173, 175, 180, 183-185, 188, 194-196, 219, 226, 255-283
 – instantaneous vs. extended 255-258, 266, 282, 283
 – simple vs. complex 265-267, 269-271, 282, 283
 – occurrent type 255, 259, 261, 267-269, 272-283
 Oehler, K. 179
 Oliver, D. 98, 101, 136
 Olson, E. T. 34
 oneness 42, 43
 ontological engineering 15, 21, 57, 65, 69, 70
 ontological abstraction (see abstraction, ontological)
 ontological sextet 188, 219-221, 226-228
 ontological square 186-188
 ontological view 84, 90

- ontology
- application ontology 171
 - applied ontology 15, 19-22, 27, 220
 - formal ontology 14, 15, 17, 22, 39-41, 43-48, 52-56, 67, 174, 196, 201, 281
 - material ontology 44, 46, 48
 - reference ontology 17, 171
 - upper-level / domain-independent 204
- OpenCyc, OpenCyc Upper Ontology 74, 173-175, 177, 196
- organism 18, 22, 23, 25, 26, 31-36, 46, 47, 92, 95, 125, 183, 191, 221, 225, 252, 304
- part** 15, 25, 27, 29, 32, 33, 36, 47, 67, 106, 113, 126, 130, 133, 138, 139, 145, 181, 183, 184, 189, 191-197, 222-228, 230-235, 237, 258, 265-267, 269-281, 304
- part_of (see relation, part_of)
- parthood (see relation, part_of)
- participation relation /
- has_participant (see relation, participation)
- particular (as a noun) 18, 34, 42, 43, 45-49, 54, 55, 85, 91, 92, 101-107, 110-113, 120, 122, 167, 173, 175, 180, 182, 185, 186, 188, 196, 200, 219, 220, 221, 226-228, 231, 240, 241, 246, 255-257, 261, 266, 267, 273-275, 278, 307
- particular (as an adjective modifying judgment, object, or similar terms) 9, 10, 12, 15, 16, 29, 45, 46, 54, 55, 65, 99, 105, 160
- partition (see granular partition)
- partition theory / theory of granular partitions (see granular partition theory)
- Peirce, C. S. 194, 255-256
- perspectivalism 23-26, 80
- physical abstraction (see abstraction, physical)
- Plato / Platonic 42, 55, 71, 159, 164, 178, 206, 207, 285
- Popper, K. 19, 285, 289, 292, 293-299, 308
- possibilia / possibility / possible 52, 54-56, 76, 77, 91, 99, 100, 140, 206, 288
- pragmatist conceptualism (see conceptualism, pragmatist)
- precise / precision (of classification system) 116, 161, 304
- predicate (verb and noun) 63, 71, 109-111, 127, 146-148, 177-180, 185, 189, 229
- principle of individuation (see individuation, principle of)
- process (sometimes as opposed to ‘occurrent’) 15, 18, 22, 25-27, 29-33, 36, 46, 68, 69, 91, 105, 111, 183-185, 194, 195, 209, 221, 226-231, 235, 240, 247-249, 255, 264, 272, 276-283, 389, 290, 304, 308
- projection relation (see relation, projection)
- proper name (see name)
- property 43, 67-69, 81, 168, 177, 186, 192, 242
- proposition 43, 70, 133, 216, 224, 232, 300, 307
- Pseudo-Dionysius 42, 55
- psychological view 84, 85, 99, 90

- quality** (property, attribute, accident, characteristic, trait) 40, 68, 111, 114, 179, 183, 189, 240, 241, 243, 244, 250
- Rea, M. C.** 203
- realism / realist** 11, 13, 14, 16, 19, 57-58, 79, 80, 90-98, 101, 103-107, 109, 111, 113-119, 123, 236, 238, 240, 244
- epistemological realism 19, 289, 292-294, 302, 308
- realist logic** (see logic, realist)
- realist fallibilism** (see fallibilism, realist)
- realist orientation** 15, 16
- representationalism** (see representationalism, realist)
- reality** 12-19, 22-25, 28, 34, 36, 41, 48-57, 62, 66, 70-81, 87, 90-98, 101, 103-107, 109, 111, 113-120, 123-128, 133, 135-138, 142, 146, 157, 158, 163, 171, 173, 184, 186-188, 197, 201, 203, 206, 212, 219, 222, 223, 226, 232, 240, 277, 285, 290, 293, 294, 296, 298-300, 302, 306, 309
- Rector, A.** 94-96
- redundant / redundancy / non-redundancy** (criterion for ontology building) 89, 104, 117, 144, 148, 149, 156, 204, 205
- refer / reference** 10, 18, 19, 27, 28, 32, 36, 43, 45, 48, 50, 51, 55, 56, 63, 75, 83, 87, 90-95, 97, 99, 100, 102-104, 106, 107, 111, 117, 120, 123, 124, 134, 165, 168, 179, 187, 189, 295, 212, 220, 221, 226, 238, 246, 259, 262, 263, 271, 274, 277-229, 282, 283, 288, 297-299, 305, 306, 308
- reference ontology** (see ontology, reference)
- reflective (judgment) / reflection / concept of reflection** 48, 52-56, 66
- regional ontology** (see ontology, regional)
- relation**
- causal 8
 - exemplification 67, 188, 219
 - inherence (inheres_in) 67, 111, 188, 219, 220
 - instance-level 214, 221, 230, 232
 - is_a 18, 30, 36, 95, 112, 113, 116, 175, 194, 220, 222-224, 226-230, 232-254, 304, 305
 - location / located_in 30, 90, 135-137, 142, 146, 147, 157, 222, 231, 232
 - parthood / part_of / part-whole 30, 36, 111, 112, 114, 116, 130, 138, 140, 144, 145, 192, 194, 220-223, 226-228, 230-233, 245, 251, 304
 - participation / participates_in 113, 188, 195, 219, 220, 222, 226, 228, 231
 - projection relation 126, 133-139, 142, 144, 145, 147, 149, 151, 153, 155, 157, 158
 - subunit relation 127-131, 139-145, 148, 150, 151, 153, 157
 - subsumption 18, 61, 175, 177, 220, 226-229, 239-252, 304, 305
 - determinable subsumption 18, 235, 241, 246, 247, 250, 251
 - genus-subsumption 18, 235, 241-244, 246, 247, 249-252
 - type-level 220-222, 225, 229, 232

- relativism 12, 77, 78
- representation / conceptualization
12-14, 18, 30, 48-52, 58, 63, 64,
70-81, 83, 100, 117, 137-144,
171, 185, 226, 233, 282, 283,
286, 293, 205, 306, 307
- context representation (a
desideratum of Cimino's) 89
- representationalism, realist 70
- reproduction / reproductive (in
context of biological
classification) 207, 209, 211,
213, 126
- Reydon, T. A. 207
- Ridder, L. 191
- Rödl, S. 267
- Rosse, C. 30
- Russell, B. 161, 183, 186
- Russell, S. 69
- Ryle, G. 264
- Schark, M. 204
- Schmidt, K. 167, 170
- Schuhmann, K. 154
- science
- computer science 61, 72, 235
- information science 7, 8, 14, 21,
22, 27, 32, 53, 69, 71, 164, 174,
235, 247, 291, 303
- natural science 19, 21, 25, 26,
34, 36, 39, 41, 43, 46, 54, 56, 60,
77, 79, 81, 91-93, 96, 97, 105,
107, 113, 117, 118, 123, 159, 160,
171-173, 207, 212, 219, 221, 222,
227, 229, 234, 238, 239, 243,
285-287, 289, 290, 292, 293, 295,
296, 299
- Searle, J. 135, 242, 307
- Sellars, W. F. 25-6
- semantic / semantics 10, 28, 29, 70,
74, 88, 97, 114, 138, 161, 166,
224, 240, 285, 293
- set theory / set-theoretic 33, 76,
126, 127, 132, 137, 138, 143,
148, 149, 189, 191-193, 215, 228,
249, 254
- Shortliffe, E. 98, 101
- Simons, P. M. 180, 191, 228
- Smith, B. 10, 11, 15-17, 21, 28-29,
32-33, 36, 43-44, 72, 80, 83-86,
94, 96-97, 109, 111, 115, 119,
120, 123, 125-127, 131, 133-5,
138, 144, 154, 156, 162, 165,
171, 174, 183, 185-186, 188, 192,
204, 219, 223, 226-227, 235, 238,
244, 281, 303
- SNAP entity (see SNAP/SPAN
ontology, continuant)
- SNAP/SPAN ontology (see also
occurrent, continuant) 185
- SNOMED / SNOMED CT 28, 73,
85, 103, 115, 166, 167, 222
- Sowa 97, 114, 174, 194-196
- Sowa Diamond 17, 174, 175,
194-196
- space 47-55, 76, 111, 133, 152-157,
190, 191, 195, 196, 213, 226,
248, 280, 281
- SPAN entity (see SNAP/SPAN
ontology, occurrent)
- Spackman, K. 98, 101, 119
- specialization 235, 248-254
- species (see also class, category,
kind, type, universal) 15, 18, 22,
23, 26, 27, 35, 36, 125, 128, 134,
139, 140, 144, 149, 155, 160,
161, 164, 169, 179, 180, 187,
193, 197-217, 241-244, 247

- biological species 33, 126, 142, 201, 204-216
- specific judgment (see judgment, specific)
- specification 18, 71, 165, 210, 202, 235, 236, 246-252, 273, 280
- state of affairs / states of affairs 48, 133, 173, 188, 189, 300, 306
- Stekeler-Weithofer, P. 54
- Stell, J. G. 134
- Stoecker, R. 23
- Stout, R. 264, 268
- Strawson, P. F. 268
- structure 4, 11, 13, 17, 24, 25, 27, 28, 30, 31, 33, 36, 41, 44, 52, 53, 58, 59, 63-70, 76, 88, 102, 103, 105, 115, 117, 124, 197-200, 204, 216, 217, 219, 221, 227, 230-232, 243, 245, 247, 250, 251, 258, 260, 262, 264, 266, 267, 269, 277, 282, 283, 301, 304
 - general / generic 257, 258, 260, 263
 - mereological (see mereology)
 - temporal / of occurrents 18, 255, 259-261, 264, 265, 267, 269, 274, 276
- Suárez, F. 40
- subjectivism / subjective 45, 75, 77, 290
- substance / substantial 31, 32, 35, 49, 67, 68, 95, 107, 133, 164, 178, 180, 189, 227, 228, 236, 249, 253
- subsume, subsumption (see relation, subsumption)
- SUMO (Suggested Upper Merged Ontology) 17, 74, 173, 174, 192
- symbolic language (see language, symbolic)
- syntax / syntactic 28, 74, 88, 110, 117, 306
- system
 - biological / biomedical 17, 29, 31, 35, 81, 114, 123, 168, 214, 256, 299
 - category 10, 14, 17, 22, 23, 26, 28, 69
 - classification/classificatory 16, 161-163, 165, 166, 169, 175, 177, 197, 207, 213, 214, 216
 - expert 61-62
 - information 7, 9-11, 13-16, 18-21, 28, 30, 31, 34, 35, 62, 69, 72, 78, 81, 89, 97, 113, 164, 165, 221
 - natural / artificial 205, 212, 214
 - terminological 70, 100
 - taxonomy / hierarchy 14, 22-26, 29, 68, 86, 155, 157, 159-162, 165, 168, 180, 193, 194, 198, 200, 204, 205, 211, 212, 214, 216, 217, 236, 242-244, 253, 255, 264, 275, 278, 283, 304, 305
 - Linnaean 31, 211
- teleology / teleological / telic / atelic 32, 262-265
- Temmerman, R. 87
- temporal entity (see occurrent)
- term 8, 14, 18, 19, 27-30, 34-36, 45, 56, 65, 69-71, 75, 76, 83-124, 160, 162, 164, 165, 168, 170, 177-179, 185, 192, 206, 221-224, 238, 240, 250, 285, 296, 300-309
- term orientation 10, 85
- terminology (controlled vocabulary) 14, 35-36, 64-66, 75, 87, 162, 222, 223, 227, 304
- Thompson, M. 208

- time 9, 10, 18, 21, 32-34, 46-55,
69, 80, 92, 93, 96, 106, 107, 111,
116, 118, 120-122, 134, 157, 177,
179, 183-185, 190, 191, 195, 196,
204, 211, 215, 219, 221-231, 237,
238, 249, 255-273, 275-278, 280-
283
- trait 9, 12-14, 32, 98, 160-166, 225
- transcendental / transcendence 41-
43, 53, 55
- transitivity / transitive 128-130,
228, 233, 236, 249
- transnatural abstraction (see
abstraction, transnatural)
- transparency / transparent 135, 136,
139, 149, 187, 300
- tree 31, 40, 132, 133, 139-142, 157,
163-166, 198-200, 216, 235, 247,
277-278, 283
- taxonomic 197, 199, 200, 215
- evolutionary 214-217
- Trettin, K. 186
- true / truth 7, 8, 40, 42-45, 51-53,
55, 78, 90, 94, 104, 112, 121,
133, 135, 141, 158, 161, 199,
214, 223, 223, 229, 232, 245,
246, 259, 273, 287-302
- truthlike / truthlikeness 289, 292-
299, 302, 308
- type (see also category, class, kind,
species, universal) 18, 23, 27, 31,
33, 36, 41, 64, 67, 89, 91, 95, 97,
102, 105, 111, 113, 121, 160,
161-168, 170, 171, 215, 220-226,
229, 132, 233, 238, 255-270, 272-
280, 282
- UMLS (Unified Medical Language
System) / UMLS Metathesaurus
28, 29, 74, 98, 99, 100, 101, 170,
222, 224
- understand / understanding 7, 9, 11-
15, 21, 22, 27, 29, 31-36, 44, 48-
49, 51-53, 56-57, 59-61, 72, 75,
77, 79, 80, 83, 89-90, 94, 98, 100,
102, 111, 114-115, 171, 173, 187,
197, 202, 208, 215, 216, 219,
224-225, 228, 268, 277, 300
- uniform / uniformity (of
classification system) 88, 92,
160, 167, 175, 183, 266
- unit (in context of granular theory)
12, 25, 85, 87, 126-131, 134-138,
140, 156, 162, 281
- universal (see also category, class,
kind, species, type) 62, 71, 91,
97, 101-102, 104-106, 109-110,
112-113, 117-121, 123, 144, 168,
177, 182, 185, 187-188, 192-193,
196, 210, 220, 222-223, 225-229,
232, 238, 240, 289, 299
- use (noun and verb) 7-8, 14-15, 17,
19, 21, 35-36, 40, 42-45, 48, 50,
52, 53, 55-56, 58, 63, 65-66, 70-
71, 80, 84-85, 94, 97-99, 101-
103, 107, 109-110, 113, 116, 118,
128, 134, 142, 146, 160, 170-171,
178-179, 192, 198, 200, 204, 206,
212, 220-225, 227-228, 238, 240,
244, 248, 253, 255, 257, 262,
270, 272, 275, 285, 290, 294,
296, 300-301, 303-306
- use / mention distinction 19, 36,
285, 300-301, 303-306
- user interface 61
- vague / vagueness / non-vagueness
(criterion for ontology building)
27, 89, 104, 117
- Vaihinger, H. 285, 297-298, 300
- Varzi, A. C. 134, 225
- vector / vector system 120

- Vendler, Z. 262-264
verbs 65, 236, 249
– transitive / intransitive 249
veridical / veridicality 11-13, 16,
80
verisimilitude 292-295
Vizenor, L. 306
vocabulary / controlled vocabulary
(see terminology)
Von Wachter, D. 186
- Williams, D. C.** 186
Wilson, J. 34
Wolff, C. 40
- WordNet 70, 301, 303
world (reality, understanding) 7-8,
10-13, 15, 19, 24-26, 32-33, 42,
43, 45, 47-48, 51-55, 57-62, 64,
66-67, 71, 78-79, 81, 83-84, 86-
88, 90-92, 95, 97-101, 103, 105,
107, 111, 113, 114, 116, 119, 121-
125, 133-135, 173, 184-185, 188,
190, 193, 203, 222, 224, 235-236,
264, 275, 280, 285, 293, 297-303,
305-306, 308
- Wüster, E. 84-88, 91, 93, 99, 101-
102, 113