

A Game of Essence and Serendipity: Superb Owls vs. Whisking Woodpeckers

Guendalina Righetti

Free University of
Bozen-Bolzano
Bolzano, Italy

guendalina.righetti@stud-inf.unibz.it

Oliver Kutz

Free University of
Bozen-Bolzano
Bolzano, Italy

oliver.kutz@unibz.it

Daniele Porello

University Of Genoa
Genova, Italy

daniele.porello@unige.it

Nicolas Troquard

Free University of
Bozen-Bolzano
Bolzano, Italy

nicolas.troquard@unibz.it

Abstract

The representation of everyday concepts is important for a number of applications, ranging from the Semantic Web to NLP and general AI. We propose here a detailed case study of the *Leuven concept database* (LCD), which is a rich database of commonsense knowledge, written in natural language. We aim to convert the commonsense knowledge contained in the LCD into a format suitable for implementation and practical application. We then investigate a hybrid approach that combines a syntactic analysis of the surface structure of the LCD entries with a semantic and ontological analysis of those entries, considering also the role of other cognitively-grounded facets of core knowledge. The approach therefore suggests a systematic portfolio of disambiguation modes with the goal of improving the match between everyday meaning of concepts and formal semantics. Finally, we illustrate the practical usefulness of this approach in a concrete computational implementation for concept combination.

Introduction

Commonsense knowledge and specifically the representation of everyday concepts is a crucial ingredient in many applications, ranging from the Semantic Web to NLP and general AI. The word “commonsense” groups different aspects of human knowledge, which permeate our experience of the world and allow us to move therein. Commonsense knowledge includes our ability to distinguish between single objects and classes of objects, to distinguish between animate and inanimate things, but also more mundane knowledge: the fact that fish live only in water (and normally do not have a job), the fact that vehicles need fuel, or the fact that my dad is necessarily born before me. Commonsense knowledge is acquired by humans through experience and throughout life in an almost completely effortless way. Despite the long tradition of research (McCarthy, 1959; Lenat, 1995) investigating how to bring this kind of knowledge from human to machine, it is still a wide-open research question. At the same time, any progress in this field directly benefits a number of AI applications.

As a case in point, in the context of Computational Creativity the representation of commonsense knowledge is crucial when dealing with Computational Conceptual Blending. In Cognitive Linguistics, Conceptual Blending has been

proposed as a general cognitive process underlying, among others, the human ability to creatively integrate and combine concepts (Boden, 1998; Fauconnier and Turner, 1998). Accordingly, a blend is constructed by selectively mapping the shared features of different (mental) input spaces into a generic, shared, mental space. The blend develops then its own emergent structure, which derives from the combination of the projected features. For humans, this process may happen imperceptibly, by exploiting information they possess, specifically relying on their commonsense knowledge. Arguably, some of the most interesting blends originate from the resolution of clashes stemming from the commonsense information which is coded into, and sometimes hidden in, the input concepts.

Computational Conceptual Blending (CCB) aims at formally interpreting and capturing the process of conceptual blending and integration. Different, though related, frameworks have been proposed in the literature, either to formally model or to replicate the process of conceptual blending (Eppe et al., 2018; Neuhaus et al., 2014; Veale, 2019; Ontanón and Plaza, 2010; Hedblom, Righetti, and Kutz, 2021; Gonçalves, Martins, and Cardoso, 2017). Rather obviously, computational systems are forced to reason with the information they are presented with, and to bootstrap the clashes and the blending process in general, CCB systems need commonsense knowledge to be represented in the input spaces. Beyond the study of the heuristics involved in the computational process of creatively blending concepts, it is then also worth focusing on the formalisation of the commonsense information which is needed as a propellant to steer the whole process.

To this end, we focus here on a detailed case study of the *Leuven concept database* (De Deyne et al., 2008; Ruts et al., 2004), which is used as a source of commonsense knowledge to be converted into a format which will be then suitable for practical application. The *Leuven concept database* (LCD) contains information, gathered by a group of psychologists at the University of Leuven, over the features exhibited by 15 category labels (here often referred to simply as *concepts*), and provides evidence on human conceptualisation. The conceptualisations that emerge from the LCD do not necessarily reflect a *good* definition of the concepts involved—at least not in a normative sense. It aims at being a good description of what people have in mind when

they think about those concepts, and of the meaning they associate with them. Therefore, the database is permeated by “commonsense information”, and exhibits some of the basic ambiguities related to the use of natural language. These conceptualisations, therefore, constitute an excellent point of observation on the challenges to be faced to make this information machine interpretable. We propose here a study which addresses exactly these difficulties. In order to make the content of the LCD available for practical application, a process of formalisation is needed: we exploit here the Web Ontology Language (OWL) as a prominent starting point. Being a computational, logic-based language, OWL obviously imposes certain limitations in terms of expressivity. The translation from the LCD to OWL thus involves a trade-off between the language’s expressive power and the desire to preserve as much information as possible. Another boundary in the translation is set by the presence of background foundational ontology distinctions, which are used to inform some of the formalisation choices in the process of translation. In particular, we argue, and present some examples, that exploiting deep ontological distinctions enables us to impose order and coherence (when possible) to the information in the LCD, helping also to disambiguate some of the hidden meaning within the data.

Finally, we conclude the paper illustrating the practical usefulness of this approach in a concrete computational implementation for concept combination. This will here serve as a demonstration and a possible use of the resulting formalised commonsense knowledge.

Related Work

Many practical AI applications require complex inferences, which, in turn, require large common-sense knowledge bases. Typical examples are chatbots, or domestic applications, e.g. involving ‘intelligent’ cooking or cleaning assistants, which need to navigate human spaces with a sufficient level of the involved common sense inferences (Krieg-Brückner et al. (2015); Bateman et al. (2018)). In practice, this need has often resulted in the use of structured lexical databases, semantic networks, or linked data, such as WordNet (Fellbaum, 2005), ConceptNet (Speer, Chin, and Havasi, 2017) and DBpedia (Auer et al., 2007) as a link between natural language and higher level semantic representations. Despite their usefulness, these repositories often show some level of ambiguity, which demonstrates the lack of a common agreement on the meaning of the lexical entries. In order to overcome this difficulty, a number of works have proposed different approaches to provide these databases deeper semantic support (Fellbaum and Hicks, 2019; Silva, Freitas, and Handschuh, 2016; Gangemi et al., 2012; Schmidt et al., 2019; Gangemi et al., 2003). The key ideas behind those approaches is to make these repository “ontology-like”, as far as possible.

In order to achieve this level of formalisation, many of the approaches mentioned above appeal to foundational ontology (FO — such as BFO, DOLCE, GFO, SUMO, etc.) which provide a common vocabulary through imposing fundamental ontological distinctions. In (Gangemi et al., 2003),

for example, a connection is drawn between WordNet’s upper level synsets and the foundational ontology DOLCE, and, more recently, (Silva, Freitas, and Handschuh, 2016) enlarged that alignment in order to include also verbs. In (Schmidt et al., 2019), a complete manual alignment between WordNet and a different Upper Ontology (SUMO) is proposed. Continuing that tradition, (Gangemi et al., 2012) propose a tool for automatically typing DBpedia entities, which relies on the alignment to both Wordnet supersenses and a subset of DOLCE Ultra Lite classes. Crucially, these works often use a top-down approach which propagates certain top level distinctions of the foundational ontology onto the more general entries in the database at hand, exploiting its internal relation (e.g. the hyponym relation).

We follow here a related but different strategy, based on a detailed case study of the Leuven Concept Database (De Deyne et al. (2008); Ruts et al. (2004)). Instead of assuming a specific FO and propagating its distinction through the database, we exploit the inverse, bottom-up, direction. We analyse the intended meaning of the information contained in the LCD and individuate seven *modes of disambiguations*, i.e. seven high level distinctions, ranging between ontological and cognitively relevant ones, which implicitly underlie the content of the LCD. Once individuated, these distinctions steer the analysis of the database, and thus the rendering choices of our translation into OWL.

We carried out the translation into OWL manually. There exists different tools for automatic natural language to OWL translation (Völker, Hitzler, and Cimiano, 2007; Emani et al., 2019; Draicchio et al., 2013; Nguyen, Razniewski, and Weikum, 2021). In order to be effective, these tools require very clear assertions and showing a regular structure. In contrast, the commonsense features collected in the Leuven concept database, in most cases, do not show this kind of regularity and lack of ambiguity that these tools presuppose.

The Leuven Concept Database

Data gathering

The Leuven concept database¹ is a large-scale data set that associates sets of features both to concepts (or categories’ labels, e.g. *Bird*) and to exemplars (or lexical entries, e.g. *magpie*). The data collection was carried out by the ConCat group at the University of Leuven from 2004 to 2008 (Ruts et al., 2004; De Deyne et al., 2008), and it consists of 15 categories and 420 associated exemplars. More precisely, the data set covers the domain of animals (**birds**, **fish**, **insects**, **mammals**, **reptiles** together with **amphibians**, with an average of 25 exemplars for each category label, and a total of 131 exemplars), and it collects information on the artifact domain (**musical instruments**, **tools**, **vehicles**, **clothing**, **kitchen utensils**, **weapons**, for a total of 169 exemplars over the six categories), on **fruit** and **vegetables** (for a total of 60 exemplars) and activities (**professions** and **sports**, again for 60 exemplars). At least a thousand students were involved in the experiments. All the material was collected in Dutch, but also an English translation is provided to make

¹Available at <https://simondedeyne.me/data>

the data available for further experimental and modelling approaches.

The studies conducted at the University of Leuven are placed in the debate between the Prototype Theory and the Exemplar Theory (Storms, De Boeck, and Ruts, 2000), and therefore present a series of experiments that aim to investigate aspects of one or the other theory. We are here mostly interested in the studies pertaining to a *feature-generation task*, where subjects were asked to provide lists of features in relation to the 15 category labels presented in **bold** above.

Participants' responses to the feature generation task were manually aggregated and adjusted with minimal stemming. Information was retained on the features' production frequency, which can be considered an indirect measure of their importance. Further, the importance of the features was also directly assessed by asking the participants to explicitly rate the importance of each feature in the definition of the concept for which they were previously generated (De Deyne et al., 2008). Figure 1 shows an example of the features generated for the category label *Bird* (see before the column MEAN). The table displays the features in Dutch and their translation to English. The numbers displayed in the table correspond to the importance ratings assigned to each feature by the participants of the experiments. The rating scale ranged from +3 (very important feature) to -3 (very unimportant feature). Globally, the feature generation task for the category label produced 28 features for the concept *Bird*.

Large-scale data-sets analysing the features exhibited by different concepts are quite rare in the literature, even though the possibility of using Amazon Mechanical Turk has made them more frequent (Vinson and Vigliocco, 2008; Buchanan, Valentine, and Maxwell, 2019). The LCD, however, shows some peculiarities that make it particularly suitable for our analysis: not only is it organised in an easily reusable shape, which makes it useful from a practical point of view; it also contains information about the importance of the features collected, which makes it interesting from a theoretical perspective, as will be explained in what follows.

The Leuven Concept-bundle

As it can be seen in the table for the concept *Bird*, the features collected relate to different aspects of birds, that range from habits ("builds nests", "eats worms"), to body parts or shape ("has wings", "has a beak"), to abilities ("can fly", "sings"), but that also pertain to more general cultural information (e.g. "is sometime kept as a pet", "is sometimes eaten by man"). A similar situation occurs for each of the concepts analysed, and if one takes a step back and looks at the features contained in the Leuven concept database as a whole, a rather fascinating picture emerges. Different features reflect different facets of the concepts involved, that in some cases barely stand together in the same description. For instance, some of the features of the concept *Fish* ("breathes under water", "has gills", "lays eggs", "lives in the sea") suggest a quite general definition of *Fish*, which relates to a somehow biological perspective on the concept. Other features describe instead the concept *Fish* in its relation with humans beings—and maybe with some subject's personal experience: some of them ("swims in aquarium", "is sometimes

kept as a pet") focus on the pet dimension of *Fish*, while others ("contains omega3", "is tasty", "sometimes smells") relate to the food dimension. Also, features pertaining to different dimensions may be considered conflicting—at least at some level: does the fish live in the sea or in the aquarium? Is it a pet or is it tasty? Similar considerations apply to all the concepts in the database: a *Sport* is a hobby, is relaxing and is fun, but can also be a *Profession*, which in turn is defined as a source of stress and frustration (but also an activity which is advantageous for the society and the economy). *Clothes* protect against the cold, but can be a status symbol, they protect from the rain but express your personality; a *Tool* is an aid, but you can injure someone with it; *Vehicles* are polluting, but they may be environmentally friendly. This gives us an idea of the context sensitivity of everyday concepts (Yeh and Barsalou, 2006), but reveals also their polysemy—the fact that those category labels are used as umbrella words for slightly different meanings. This also recalls what Hofstadter (2001) called the process of chunking: the idea that humans build their concepts by gluing several concepts together through their lifetime, so that at the end a concept results in "nothing but a tightly packaged bundle of analogies".

Taking a step forward and looking at the features more closely in terms of formalisation possibilities, some problems quickly emerge, e.g. regarding precision. One of the problems may be summarised as a lack of implicit knowledge. This does not only refer to the lack of fundamental categorical distinctions (see below), but also to the omission of some of the things that subjects may have considered obvious during the experiment. Subjects tend to omit some of the most obvious features (e.g. that a fish "has two eyes"), trying to focus on the more distinguishing ones (De Deyne et al., 2008). Also, they fail to specify some underlying knowledge, which they may consider not necessary for general understanding—fish are said to "swim in aquarium", compressing the more detailed information "some fishes swim in water contained in some aquarium". Another problem is the presence of errors within the data, most of which correspond to a naïve use of the "is-a" relation: a *Fish* is said to be a shark, a *Tool* is a hammer, etc.

From the LCD to OWL Ontologies

Interpreting the features

The problems described in the section above would not cause any issue for human understanding, which shows great flexibility in interpreting natural language sentences. However, when the goal is to make the features machine interpretable, they require some adjustments. Let us, for instance, consider to translate 'naïvely' the feature "swims in aquarium" into an OWL axiom, and to add it to an ontology of *Fish*. In the ontology there could be a definition of *swims*, maybe as an action which is performed only in a particular environment—namely in water². In order to avoid inconsistencies, such as the identification of 'aquarium' and 'water', one may need then to fully specify the meaning and function

²Unless one wants to consider a metaphorical use of the word *swim*, which would make the situation even more complicated.

DUTCH	ENGLISH		...		MEAN	PF	Syntax	ModesOfDisambiguation	Possible OWL Rendering		
heeft veren	has feathers	3	3	...	3	3	2,833	20	Default	Mereology	Bird SubClassOf hasBodyPart some Feather
heeft vleugels	has wings	3	3	...	3	3	2,75	9	Default	Mereology	Bird SubClassOf hasBodyPart some Wing
kan vliegen	can fly	2	3	...	3	3	2,75	20	Default/Exist.	Ability	Bird SubClassOf hasAbility some FlightAbility
heeft een snavel	has a bill	3	3	...	3	1	2,75	15	Default	Mereology	Bird SubClassOf hasBodyPart some Bill
bouwt nesten	builds nests	3	3	...	3	3	2,66	16	Default	Ability	Bird SubClassOf hasAbility some BuildingNest
heeft twee vleugels	has two wings	3	3	...	3	2	2,66	3	Default	Mereology	Bird SubClassOf hasBodyPart exactly 2 Wing
legt eieren	lays eggs	3	3	...	3	2	2,583	20	Default	Ability	Bird SubClassOf hasAbility some ReproductionAbility
heeft een bek	has a beak	3	3	...	3	2	2,5	6	Default	Mereology	Bird SubClassOf hasBodyPart some Beak
heeft twee poten	has two paws	3	3	...	3	1	2,5	3	Default	Mereology	Bird SubClassOf hasBodyPart some Paw
is een dier	is an animal	3	3	...	3	1	2,416	3	Universal	Rigid	Bird SubClassOf Animal
fladdert	flutters	2	2	...	-1	3	2,166	2	Default	Ability	Bird SubClassOf hasAbility some FlutterAbility
eet wormen	eats worms	3	2	...	2	1	1,66	8	Default	Image schema (Containm.)	Bird SubClassOf eats some Worm
fluit	sings (whistles)	2	2	...	2	3	1,333	7	Default	Ability	Bird SubClassOf hasAbility some WhistleAbility
tsjilpt	chirps	1	2	...	2	2	1,333	6	Default	Ability	Bird SubClassOf hasAbility some ChirpAbility
eet kleine dieren	eats small animals	2	1	...	1	0	1,333	2	Default	Image schema (Containm.)	Bird SubClassOf eats some SmallAnimal
leeft in het wild	lives in the wild	2	2	...	3	1	1,083	3	Default	SpatioTemporal	Bird SubClassOf hasLocation some WildArea
is een trekvogel	is a migratory bird	2	1	...	-1	-1	1	4	Default	Anti-rigid	MigratoryBird SubClassOf Bird
vind je in bomen	can be found in trees	1	2	...	-2	-1	1	5	Existential	SpatioTemporal	Bird SubClassOf hasLocation some TreeArea
heeft poten	has legs	0	1	...	1	2	0,75	3	Default	Mereology	Bird SubClassOf hasBodyPart some Leg
in kooi	lives in a cage	1	2	...	1	-1	0,666	2	Default	SpatioTemporal/Image Sche	Bird SubClassOf hasLocation some CageLocation
heeft luchtzakken	has air sacs	-3	2	...	2	1	0,666	3	Universal	Mereology (Essential part)	Bird SubClassOf hasBodyPart some AirSac
kraaloojjes	has beady eyes	3	3	...	2	-2	0,5	2	Default	Mereology	Bird SubClassOf hasBodyPart some BeadyEye
zingt	sings	1	1	...	1	2	0,416	3	Default	Ability	Bird SubClassOf hasAbility some SingAbility
kan lopen	can walk	-1	1	...	1	1	0	2	Default	Ability	Bird SubClassOf hasAbility some WalkAbility
kan zich voortplanten	is able to reproduce	-3	1	...	-3	2	-0,5	4	excluded	excluded	excluded
wordt soms als huisdier	is sometimes kept as a pet	-3	0	...	-2	-2	-1,16	2	excluded	excluded	excluded
wordt soms door mensen gegeten	is sometimes eaten by humans	-3	0	...	-3	-3	-1,25	2	excluded	excluded	excluded

Figure 1: **Bird**: an example of a “Feature via Category Label” table, plus annotations from our analysis. MEAN refers to the mean of the importance associated by the subjects to the features. PF is the production frequency.

of the object ‘aquarium’, and the image-schematic relation of containment it bears with the water.

These considerations suggest the need of a preprocessing phase, which we conducted at different levels. At the very first level, this preprocessing phase was carried out following the Gricean communication principle, or *Cooperative Principle*. In the context of any language exchange, the principle prescribes to “make your contribution such as is required (...) by the accepted purpose or direction of the talk exchange in which you are engaged” (Grice, 12 Dec 1975). In this specific case, the *purpose of the talk exchange* was the definition of the concepts proposed by the psychologists, and the *talk exchange* was more precisely the experiment completed by the participants. According to Grice, the violations of this principle, which here includes errors and imprecisions, should be interpreted in such a way as to protect the rationality of the speaker, according to Quine’s *Principle of Charity* (Quine, 1960), which prescribes interpreting a speaker’s statements in the most rational way possible, and considering its best, strongest possible interpretation³.

Despite these premises, some of the features produced in the Leuven experiments were difficult to interpret in the context of the definition of the concept—sometimes because blatantly false when stated for the whole class, sometimes because they were related to a semantic context completely different to the one proposed in the experiment. To give a taste: a *Fish* is “a constellation” and *Weapons* are “used in sport”. Some of the features, then, captured biases of our language (and our society): a *Profession* “is different for

men and women”, and a *Kitchen Utensil* “is especially used by women”, but a *Tool* “is primarily used by men”. As described in more detail above, all the features generated in the experiments were afterwards judged in order to evaluate their applicability to the class at hand (see again the numbers in Figure 1). Inspired by the prototype-theoretic notion of *salience* (Rosch, 1973), and in order to exclude some of the most controversial features, we calculated the mean of subjects’ judgments, and excluded the entries strictly below the threshold 0. This procedure allowed us to exclude 102 features, namely around 20% of the features.

The formalisation step

After the preprocessing step, the features are translated into OWL axioms. The Web Ontology Language (OWL) is one of the most widespread language for authoring ontologies. It allows the users to write explicit and formal conceptualisations of a domain model. We will just sketch here the features of the language, the interested reader may refer to (Antoniou and van Harmelen, 2009; McGuinness and van Harmelen, 2004) for a more in depth description.

In particular, OWL is a logic-based language: it is mapped to Description Logics, i.e. decidable fragments of first-order logic. This provides OWL with a clear, well defined, formal semantics and efficient reasoning services. The reasoning support is important not only to compute ontologies’ implicit knowledge (i.e. the entailed statements), and thus to reason over the axioms, but also to check their consistency, the presence of unintended consequences, etc. At the same time, efficient reasoning services require some limitations in the expressiveness of the language. Some trade-off is then necessary between the performance of the reasoning and the

³Interestingly enough, Quine developed this principle in the context of language translation.

language’s expressive power, which should allow the user to express large volumes of knowledge.

More precisely, OWL allows to express knowledge about classes, instances and binary relations between instances. It provides different constructs to declare the different entities of the language: here we mainly deal with the constructs **class**, **object property** and **individual**. A class defines a set of individuals that share some properties; **object properties** are used instead to assert binary relationships between individuals; **individuals** are instances of the classes. For example, we may want to declare the class of *Bird* as the set of those instances that share the features described above. All the 15 categories described in the Leuven concept database are indeed examples of **classes**. If we want to populate the class, we may declare Tweety as an **individual** of the class *Bird*. Consider instead the feature “builds nest”: the word ‘builds’ should be interpreted as an **object property**, which relates the instances of the class *Bird* and the instances of the class *Nest*. At the same time, the set of all entities that build nests provides another example of a class, which *Bird* is a subclass of. Classes can indeed be organised in hierarchies, according to their generality, by means of the “subClassOf” relation, which behaves like the subsumption relation in Description Logic. We may also declare that two classes are “disjoint”, having no common instance, and that two classes are “equivalent”, having exactly the same instances.

The semantics of the “subClassOf” relation implies that all the elements of the sub-class are also elements of the super-class (it is indeed the subset relation). Asserting that *Bird* is a subclass of the class of entities which *build Nest*, means then that all the instances in the class of Birds build nests, without exceptions. Obviously this is a quite strong requirement when we are dealing with natural language formalisation and everyday concepts. Some of the features are described by people by means of expressions which emphasise their partial applicability to the class into consideration (e.g. sometimes, can have, etc). In other cases, this is implicit in the use of everyday language (e.g. people may assert that birds can fly, but this does not imply that they believe that a penguin is not a bird). Also for this reason there has been some work recently trying to allow a more cognitively grounded modelling (Porello et al., 2019; Righetti et al., 2019, 2021a), as well as defeasible subsumption (Britz and Varzinczak, 2017; Casini and Straccia, 2010), which allows to handle exceptions and counterexamples.

Following these intuitions, the features collected in the LCD can be grouped in different meta-categories, according to their grammar. This classification can be thought of in terms of Aristotle’s famous *square of opposition*. We can distinguish between: i) Universal affirmative statements, i.e. the (positive) features that apply to the whole class under consideration. As an example, we may consider the statements “a Fish is an Animal”, or “a Kitchen Utensil is a Tool”. Those statements can be treated as simple class inclusion, and in First Order Logic would correspond to universal quantification (“all fish are animal”, etc). ii) Existential statements, which apply only to some instances of the class at hand: e.g. *Insect* “can bite”, or *Tool* “can be automatized”. In First Order Logic they would correspond to

existential quantification: some insect bites, some tool is automatized, etc. iii) Universal negations, which apply again to the whole category, but which express a negated statement, like *Fish* “does not live on land” or *Insect* “does not live long”. iv) Existential negation, of the kind “some A are not B”, and which apply only to a subclass of the concept under consideration: e.g. a *Vegetable* “is not always green”.

Table summarises the distribution of the features in the different meta-categories. As it can be seen in the table, most of the features enter the meta-category “Universal affirmative”, while the negated statements (Universal negative and Existential negative) are very few.

Type of Statement	Frequency
Universal affirmative	≈ 82%
Existential affirmative	≈ 16%
Universal negative	≈ 1%
Existential negative	≈ 0, 5%

Table 1: Features classification

Looking beyond the syntactic surface, however, within the Universal affirmative statements, only a few (less than 10%) are true, clear universal statements, which are valid for the whole category. Many other features (see e.g. the column ‘Syntax’ in Table 1) look like universal statements, but presuppose the possibility of exceptions: e.g. “Birds eat worms” is used as a *default* statement about birds, but it is possible to think of counterexamples, since not all birds are carnivores. When translating the features into axioms, it is desirable to distinguish between the axioms which require a classical, non-defeasible, use of the SubClassOf relation (e.g. *Bird SubClassOf Animal*), and axioms which do require a defeasible semantics (e.g. *Bird SubClassOf eats some Worm*). This distinction is registered at the level of annotation, and can be guided in different ways. In part, it is guided by the information in the database: we can in fact use the features’ production frequency and their average judgments to take some decisions. The features which are generated often and which get a high average rating are more likely to be valid for the whole category. However, this strategy alone does not always guarantee satisfactory results. The feature “has feathers”, for instance, was produced for the concept *Bird* by all the subjects involved in the experiment, and got the highest rating. However, it would be reasonable to make it a defeasible axiom, since e.g. many pullets do not have feathers.

A Game of Disambiguation

Foundational or upper ontologies (FO) formalise the meaning of very general terms, such as object, event, property, quality, relation, process, etc. (Borgo, Galton, and Kutz, 2022). They provide the top-level categories that are in principle common to many domains of application, and are implicitly at work in common sense. There are a number of different such ontologies which reflect different philosophical views on reality, ranging from a realistic stance endorsed by BFO (Arp, Smith, and Spear, 2015) to a cognitivist per-

spective enabled by DOLCE (Masolo et al., 2002). While we do not take a position here about which is the right FO to analyse commonsense concepts, we stress that embracing the perspective of a selected FO has important consequences on the formal rendering of the commonsense expressions. For the sake of this discussion and for highlighting the use of FOs in general in representing commonsense concepts, we exemplify how a number of features in the Leuven concept database can be construed by means of a foundational analysis. Despite the disambiguation choices we propose here, some of the features in the database were still too idiosyncratic to fit modelling and logical rendering strategies, and were therefore manually discarded.

We here combine a two-level approach. Firstly, we identify a candidate categorical statement elicited from the LCD (e.g. All *As* are *Bs*). Secondly, we use FOs and their distinctions to help in identifying the intended meaning of classes *A* and *B* and in understanding the relevant representational choices. Although this section is descriptive in nature, it provides the basic rules of a **game of disambiguation** governed by foundational choices and representational modes. We therefore organise the discussion along **7 basic modes of disambiguation**:

Mode 1: Rigidity and anti-rigidity Two important general properties of classes are *rigidity* and *anti-rigidity*, cf. Guarino and Welty (2004). A *rigid* class is such that every instance of that class is *necessarily* an instance of that class. For example, in Figure 1, the feature *Animal* can be intended as a rigid class: a bird is an animal and, throughout its life, cannot cease to be an animal. An *anti-rigid* class is a class such that its instances eventually cease to be instances of that class. For example, in Figure 1 we have the feature *Migratory*. This class can be interpreted as an anti-rigid class, a *phase* of the life of the birds which has a beginning and an end. So when we represent the Leuven entries by means of axioms such as “all birds are animals” and “some birds are migratory”, we can refine the meanings of these two statements by categorising the features as rigid or anti-rigid⁴. The *rigidity* and *anti-rigidity* distinction then plays a role in the context of Universal vs Existential statements described in the previous section. We may have statements of the kind “all *As* are *Bs*”, where *B* is a rigid property, which means all *As* are always *Bs* (e.g. all Bird are always Animal). But we may also have Existentials of the kind “some *As* are *Bs*”, where *B* is a rigid property, which again means that “some *As* are *always Bs*” (e.g. “some Animals are Birds”). On the other hand, both Existentials and Universals can involve anti-rigid properties, e.g. “some Birds are Migratory” and “all Mammal (mothers) breastfeed its babies”. These issues are, of course, closely related to the semantic complexities found in Aristotle’s modal Syllogistic (Malink, 2013).

Mode 2: Mereology An important ontological aspect is mereology, the theory of part-whole relations. FOs usually contain an axiomatisation of mereology, which makes the

⁴Handling the distinction between rigidity and anti-rigidity for OWL formalisations is challenging due to the limited expressive power of DLs.

meaning of parthood relations explicit. Although the parthood relation may be not overly manifested in the syntax of the description of a feature, a number of entries in the LCD contain statements about the parts of an entity. E.g. in Table 1, “have wings” clearly indicates a parthood relation. So the rendering of that statement may be an axiom that states that the class of birds is included in the class of things that are related, via the parthood relation, to the class of wings.⁵ The parthood relation is widespread in almost all the domains of the LCD (with the exception of the concept “Profession”), and constitutes about 13% of the features. The general ontological notion of parthood is quite abstract and, in many cases, one of the specialised parthood relations has to be considered, e.g. functional parthood, necessary parthood, temporary parthood, etc.

Mode 3: Spatio-temporal relations / Image Schemas

Many entries in the database (15%) specify possible places in which an entity can dwell, e.g. “lives in the wild”, “found in trees” (Bird), but also “sold in clothes shop” (Clothes) or “is often found in action movies” (Weapon). For these cases, FOs usually reify spatial and temporal locations as particulars of the ontology and can express the fact that an entity is located at a certain place or time. For instance, we can introduce the class of entities “located in the wild”. These classes can be analysed according to the rigidity vs anti-rigidity distinction that we introduced earlier to assess the strength of the constrain: it may be necessary for fish to live in the water, while only accidental (non-rigid) to live in a cage or in the wild. Particularly salient spatio-temporal relations, and also prevalent in the LCD, are image-schematic ones, such as *containment*, *support*, or *path-following*. The importance of image schemas in computational blending has been illustrated in detail by Hedblom, Kutz, and Neuhaus (2016).

Mode 4: Quality and quality spaces A number of entries refer to qualities—i.e. colours, shapes, sizes, weights, etc.—of the instances. Around 12% of the features could be understood as qualities. For these cases, FOs like DOLCE provide a quite sophisticated analysis of quality ascriptions, relying on Conceptual Spaces (Gärdenfors, 2000). This approach renders the ascription of colors by introducing a relation of *location* between a *quality* and its *quale*, e.g. between the colour of a fish and a particular value of it, such as “bluish grey”, which belongs to a suitable conceptual space of colours.

Mode 5: Constitution Other entries in the LCD contain the expression “made of” which is usually associated to what ontologists term *constitution*. For instance, DOLCE has a well-developed theory of *constitution* that is capable of approaching classical philosophical puzzles involving the persistence conditions of a statue constituted by a lump of clay. In this context, all the claims about the constitution of objects pertain to the artifact domain: *Vehicles* are “made of Metal”, *Clothes* are “made of Textile”, etc. Around 5% of the features of the artefact domain fall in this category ($\approx 2, 5\%$ when considering the whole set of features).

⁵Another technical issue is to specify that the bird has to have exactly two wings, but that this assumption is defeasible.

Mode 6: Action and ability There are entries (mostly in the domain of animals, which constitute around 16% of the features) that ascribe an ability to an *agentive object*. Agentivity is sometimes intended in a broad sense, including animals. In entries such as “swim”, “flutters”, “sings”, the intended meaning is that the animal can perform certain types of actions, i.e. the proper ontological category to assign is that of *ability*. Other entries include the word *can* which is quite challenging due to at least two meanings of *can* which have been intensively studied in knowledge representation: *ability can* (e.g. “Birds can fly”) and *opportunity can* (e.g. “Birds can be found in trees”), which does not indicate an ability of the bird).

Mode 7: Functionality and affordances Other general concepts may be found in applications of foundational ontologies, in mid-level ontologies, or in more specific domain ontologies. Deciding how general a concept is may be a matter of discussion, however we can indicate a few quite general concepts that have applications in representing Leuven entities. *Functionality* is a concept usually related to artifact ontologies. Functionalities are intended to represent the purpose or the use of an artifact. Functionality is suitable to represent features that are expressed by means of words like “use”, such as in “it is used to prepare Food” (Kitchen Utensils), or “means”, as in “is a faster means of transportation” (Vehicle). *Affordances* (Turvey, 1992) are related to functionality in that they suggest a possible use of the object involved. For instance, a “Kitchen Utensil can be used to cut things”, i.e. affords cutting. Around 20% of the features entered this mode of disambiguation.

In Figure 1, the column ‘Modes Of Disambiguation’ shows an example of the application of this analysis to the concept **Bird**. Overall, around 15% of the features contained in the database escaped our classification according to these modes of disambiguation, most of which where from the domain of activities (Professions and Sports).

Exploiting Commonsense Knowledge: The example of concept hybridisation

We conclude the paper discussing an application of the resulting formalised commonsense knowledge in a concrete computational implementation for concept combination.

To this end, we here briefly discuss the approach of Righetti et al. (2021b), who recently proposed an algorithmic modelling of the process of concept combination, leveraging the refinement operators described in Confalonieri et al. (2020) (restricted to description logic *ALC*), and the techniques of axiom weakening. The paper aims to imitate the process of making sense of ‘impossible’ hybrid combinations, i.e. combinations of clashing concepts into imaginary objects such as “a Vehicle that is a Fish”. This is inspired by the empirical research in cognitive psychology identifying human heuristics for combining concepts that lack any obvious similarities (Hampton, 2017).

In the approach of Righetti et al. (2021b), concepts are represented as formal ontologies in Description Logic, and the combination process is, thus, rendered as an ontology in-

tegration task. Briefly, the authors propose a turn-based algorithm which is initiated with two input ontologies which need to be blended into a final ontology describing the combined concept. The authors tested the procedure on the combination of the concepts Fish and Vehicle to try to replicate one of the human combinations studied in the experiments of Hampton (2017)—namely, the Fish-Vehicle concept. In this case, the concepts of interest to be combined are not just dissimilar, but, when formalised in a logical language, jointly contain obvious and sometimes hard to resolve formal inconsistencies. When adding an axiom to the combination causes an inconsistency, the approach of axiom weakening is applied until a jointly consistent compromise is found. Intuitively, a general concept inclusion axiom of the form $C \sqsubseteq D$ can be weakened by either specialising the concept C to a smaller class, or by generalising the concept D to a larger class, w.r.t. a given reference ontology⁶.

In order to replicate human concepts one needs some repository of commonsense knowledge as a source for the input ontologies. A straightforward, encyclopedic definition of the concept at hand, would hardly faithfully represent what people have in mind when blending concepts. Consider, for instance, the Fish definition from Wikipedia: “Fish are aquatic, craniate, gill-bearing animals that lack limbs with digits. Included in this definition are the living hagfish, lampreys, and cartilaginous and bony fish as well as various extinct related groups.”⁷ This definition might quite easily be axiomatised in a formal ontology, and there exist different tools for automatic natural language to OWL translation (see the related work section above) which can be employed in the presence of such clear and precise definitions. However, when combining the concepts Fish and Vehicle, people consider much more mundane knowledge. For instance, when combining the two concepts, humans may notice that while a Fish eats Food (to stay alive), a Vehicle needs Fuel (to move) (Hampton, 2017). By exploiting a heuristic similar to the analogical mapping described by Fauconnier (1997), people would tend to *generalise* this information into “both Fish and Vehicle need some kind of Energy to move”, thus creating an interesting analogy between Food and Fuel, which would further support the integration of the two concepts into the combination “Fish-Vehicle”.

Fortunately, this is exactly the kind of commonsense information the Leuven concept database is permeated with, thus suggesting the concrete usefulness of a formalisation of the concepts contained in the database for practical AI applications. To give a further tangible example, we fed the implementation proposed by Righetti et al. (2021b)⁸ with two concepts contained in the Leuven Database, namely Bird and Kitchen Utensil, previously formalised in OWL exploiting the disambiguation steps described in this paper. The concept of Kitchen Utensil-Bird was also one of the exam-

⁶We refer to the work of Righetti et al. (2021b) for the full details about the use of axiom weakening in this context.

⁷<https://en.wikipedia.org/wiki/Fish>.

⁸Available at <https://bitbucket.org/troquard/ontologyutils/src/master/>.

ples exploited by Hampton (2017) in his experiments on impossible combinations—see Figure 3.

Besides the features contained in the database, we included in the ontologies a few additional axioms, aiming at replicating some of the commonsense distinctions needed to reason about the concepts at issue but not explicitly mentioned by the subjects during the Leuven experiments because considered obvious or out of scope (as discussed above). We added, for example, the information that Animal and Tool are disjoint classes, or that if something is located in the kitchen it cannot (not normally, at least) be located on a tree, etc. An excerpt of one of the resulting ontologies for the concept KitchenUtensil-Bird is shown in Figure 2.

- hasPart **some** Beak
- eats **some** SmallAnimal
- hasPart **some** Bill
- hasPart **some** Wings
- hasLocation **only** Location
- hasPart **only**
(ArtefactPart **or** BodyPart)
- madeOf **only** Material
- hasFunction **only** Function
- hasFunction **some** ForCooking
- hasAbility **some** StandingHeat
- hasFunction **some** SimplifyWork
- hasFunction **some** BeingAWeddingPresent
- requires **some** WashingActivity
- madeOf **some**
(Metal **or** Plastic)
- hasPart **some** ElectronicComponent
- hasLocation **some** KitchenArea

Figure 2: A Bird which is also a Kitchen Utensil: an example of a blend exploiting the LCD information.

The procedure described by Righetti et al. (2021b) allows for a fine-grained selection of the combination strategies, by allowing the choice of a preference order over the axioms assigned to both agents/ontologies, as well as the distribution of turns. Also, different evaluation strategies are proposed to evaluate the outcome of the combination. Here, the example has just an illustrative purpose, and we set a random order over the axioms and an equal distribution of turns. However, the output of the procedure is surprisingly similar to the combination of the two concepts as observed in the experiments described by Hampton (2017), an example of which is presented in Figure 3.

As the output of our procedure, the *Whisking Woodpecker* has a beak and wings, thus showing body parts, but it also has artefact parts (the whisk), it is used for cooking and, being unhygienic, it also requires some washing.

Discussion and Future Perspective

The analysis of the Leuven concept database has clearly shown the existence of a mismatch between the syntactic surface form and the content (or the *intended meaning*) of people’s statements. On the one hand, we demonstrated this mismatch in the context of universal statements, where people often adopt a default reasoning strategy, and where the meaning they intend to convey is more likely close to an



“a woodpecker has been trained to whisk eggs using its powerful head movements (...) it would not need electrical power (good for camping trips) but would on the other hand be unhygienic.”

Figure 3: The *Whisking Woodpecker* (Redrawn illustration as given by Hampton (2017), page 113): A woodpecker used to whisk as imagined by one of the participants of Hampton’s experiments.

existential interpretation. On the other hand, many of the entries in the Leuven concept database lack a syntactic trigger that would help to identify their intended meaning, and ontological analysis is required to instead find a semantic trigger. So, for instance, the recognition of *mereological* assertions was mostly guided by our general language comprehension and competence for world understanding, but it was not manifested in the syntactical description of the features. Ontological analysis offers the means to explore systematically the possible meanings of commonsense feature ascriptions and, as a result, supports a more faithful formalisation into OWL.

We have also highlighted the fruitful connection between commonsense knowledge extraction and computational conceptual blending. Specifically, we have illustrated the practicality of this connection in a concrete computational workflow and use-case. Although illustrative in purpose, the example showed the effectiveness of the proposed methodology in replicating human conceptual combination as observed in the context of experimental psychology. In this context, the modes of disambiguation, as well as the syntactic analysis described above, could be exploited further, as a way to steer and guide the blending process. One may, for instance, integrate the dialogue implementation discussed above (Righetti et al., 2021b) to take into account such information, e.g. through the application of appropriate preference orders over the axioms of the two agents. This way one may prefer axioms involving universal, rigid statements, thus preserving them during the combination, and instead prefer the weakening of default statements, thus simulating a process similar to defeasible inference. Therefore, we plan to develop this further in future work towards more fine-grained evaluation metrics for blends and their creativity, in which essence fights serendipity.

Author Contributions

G. R. was the author mainly responsible for analysing the data. Aside from this, all authors equally contributed to the research presented in this paper.

Acknowledgments

This research is partially supported by Italian National Research Project PRIN2020 2020SSKZ7R.

References

- Antoniou, G., and van Harmelen, F. 2009. Web ontology language: OWL. In Staab, S., and Studer, R., eds., *Handbook on Ontologies*, International Handbooks on Information Systems. Springer. 91–110.
- Arp, R.; Smith, B.; and Spear, A. D. 2015. *Building Ontologies with Basic Formal Ontology*. Mit Press.
- Auer, S.; Bizer, C.; Kobilarov, G.; Lehmann, J.; Cyganiak, R.; and Ives, Z. G. 2007. Dbpedia: A nucleus for a web of open data. In Aberer, K., and al., eds., *The Semantic Web, 6th International Semantic Web Conference, 2nd Asian Semantic Web Conference, ISWC 2007 + ASWC 2007*, volume 4825 of *Lecture Notes in Computer Science*, 722–735. Springer.
- Bateman, J.; Beetz, M.; Beßler, D.; Bozcuoğlu, A. K.; and Pomarlan, M. 2018. Heterogeneous Ontologies and Hybrid Reasoning for Service Robotics: The EASE Framework. In Ollero, A.; Sanfeliu, A.; Montano, L.; Lau, N.; and Cardeira, C., eds., *ROBOT 2017: Third Iberian Robotics Conference*, 417–428. Cham: Springer International Publishing.
- Boden, M. A. 1998. Creativity and artificial intelligence. *Artificial intelligence* 103(1-2):347–356.
- Borgo, S.; Galton, A.; and Kutz, O. 2022. Foundational Ontologies in Action: Understanding foundational ontology through examples. *Applied Ontology* 17(1).
- Britz, K., and Varzinczak, I. J. 2017. Context-based defeasible subsumption for dsroiq. In Gordon, A. S.; Miller, R.; and Turán, G., eds., *Proc. of the 13th International Symposium on Commonsense Reasoning, COMMONSENSE 2017*, volume 2052 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Buchanan, E. M.; Valentine, K. D.; and Maxwell, N. P. 2019. English semantic feature production norms: An extended database of 4436 concepts. *Behavior Research Methods* 51(4):1849–1863.
- Casini, G., and Straccia, U. 2010. Rational closure for defeasible description logics. In Janhunen, T., and Niemelä, I., eds., *Logics in Artificial Intelligence - 12th European Conference, JELIA 2010*, volume 6341 of *Lecture Notes in Computer Science*, 77–90. Springer.
- Confalonieri, R.; Galliani, P.; Kutz, O.; Porello, D.; Righetti, G.; and Troquard, N. 2020. Towards even more irresistible axiom weakening. In Borgwardt, S., and Meyer, T., eds., *Proc. of the 33rd International Workshop on Description Logics (DL 2020)*, volume 2663 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- De Deyne, S.; Verheyen, S.; Ameel, E.; Vanpaemel, W.; Dry, M. J.; Voorspoels, W.; and Storms, G. 2008. Exemplar by feature applicability matrices and other dutch normative data for semantic concepts. *Behavior Research Methods* 40(4):1030–1048.
- Draicchio, F.; Gangemi, A.; Presutti, V.; and Nuzzolese, A. G. 2013. FRED: from natural language text to RDF and OWL in one click. In Cimiano, P.; Fernández, M.; López, V.; Schlobach, S.; and Völker, J., eds., *The Semantic Web: ESWC 2013 Satellite Events, Revised Selected Papers*, volume 7955 of *Lecture Notes in Computer Science*, 263–267. Springer.
- Emani, C. K.; Silva, C. F. D.; Fiés, B.; and Ghodous, P. 2019. NALDO: from natural language definitions to OWL expressions. *Data Knowl. Eng.* 122:130–141.
- Eppe, M.; Maclean, E.; Confalonieri, R.; Kutz, O.; Schorlemmer, M.; Plaza, E.; and Kühnberger, K.-U. 2018. A computational framework for conceptual blending. *Artificial Intelligence* 256:105–129.
- Fauconnier, G., and Turner, M. 1998. Conceptual integration networks. *Cognitive science* 22(2):133–187.
- Fauconnier, G. 1997. *Mappings in thought and language*. Cambridge University Press.
- Fellbaum, C., and Hicks, A. 2019. When WordNet Met Ontology. In Borgo, S.; Ferrario, R.; Masolo, C.; and Vieu, L., eds., *Ontology Makes Sense - Essays in honor of Nicola Guarino*, volume 316 of *Frontiers in Artificial Intelligence and Applications*, 136–151. IOS Press.
- Fellbaum, C. 2005. WordNet and wordnets. In Brown, K., ed., *Encyclopedia of Language and Linguistics*, 665–670. Oxford: Elsevier.
- Gangemi, A.; Guarino, N.; Masolo, C.; and Oltramari, A. 2003. Sweetening WordNet with DOLCE. *AI Magazine* 24(3):13.
- Gangemi, A.; Nuzzolese, A. G.; Presutti, V.; Draicchio, F.; Musetti, A.; and Ciancarini, P. 2012. Automatic Typing of DBpedia Entities. In Cudré-Mauroux, P., and al., eds., *The Semantic Web – ISWC 2012*, 65–81. Berlin, Heidelberg: Springer.
- Gärdenfors, P. 2000. *Conceptual Spaces: The Geometry of Thought*. MIT Press.
- Gonçalves, J.; Martins, P.; and Cardoso, A. 2017. Blend City, BlendVille. In *ICCC-17*, 112–119.
- Grice, H. P. 12 Dec. 1975. *Logic and Conversation*. Leiden, The Netherlands: Brill. 41 – 58.
- Guarino, N., and Welty, C. A. 2004. An Overview of OntoClean. In *Handbook on Ontologies*. Springer. 151–171.
- Hampton, J. A. 2017. Compositionality and concepts. In Hampton, J. A., and Winter, Y., eds., *Compositionality and Concepts in Linguistics and Psychology*. Cham: Springer International Publishing. 95–121.
- Hedblom, M. M.; Kutz, O.; and Neuhaus, F. 2016. Image schemas in computational conceptual blending. *Cognitive Systems Research* 39:42–57.
- Hedblom, M. M.; Righetti, G.; and Kutz, O. 2021. Deciphering The Cookie Monster: A Case Study in Impossible Combinations. In de Silva Garza, A. G.; Veale, T.; Aguilar, W.; and Pérez y Pérez, R., eds., *Proceedings of the Twelfth International Conference on Computational Creativity*, 222–225. Association for Computational Creativity (ACC).

- Hofstadter, D. R. 2001. *Epilogue: Analogy as the core of cognition*. Cambridge, MA: MIT Press. 499 – 538.
- Krieg-Brückner, B.; Autexier, S.; Rink, M.; and Ghomsi Nokam, S. 2015. Formal Modelling for Cooking Assistance. In De Nicola, R., and Hennicker, R., eds., *Software, Services, and Systems*, volume 8950 of *Lecture Notes in Computer Science*. Springer International Publishing. 355–376.
- Lenat, D. B. 1995. CYC: A Large-Scale Investment in Knowledge Infrastructure. *Commun. ACM* 38(11):33–38.
- Malink, M. 2013. *Aristotle's Modal Syllogistic*. Harvard University Press.
- Masolo, C.; Borgo, S.; Gangemi, A.; Guarino, N.; Oltramari, A.; and Schneider, L. 2002. WonderWeb deliverable D17. The wonderWeb Library of Foundational Ontologies and the DOLCE ontology.
- McCarthy, J. 1959. Programs with Common Sense. In *Proceedings of the Teddington Conference on the Mechanization of Thought Processes*, 75–91. London: Her Majesty's Stationary Office.
- McGuinness, D. L., and van Harmelen, F. 2004. OWL Web Ontology Language Overview. Technical report, W3C Recommendation.
- Neuhaus, F.; Kutz, O.; Codescu, M.; and Mossakowski, T. 2014. Fabricating monsters is hard: towards the automation of conceptual blending. In *Proc. of the 3rd Int. Workshop on Computational Creativity, Concept Invention, and General Intelligence (C3GI@ECAI-14)*.
- Nguyen, T.; Razniewski, S.; and Weikum, G. 2021. Advanced semantics for commonsense knowledge extraction. In Leskovec, J.; Grobelnik, M.; Najork, M.; Tang, J.; and Zia, L., eds., *WWW '21: The Web Conference 2021, Virtual Event / Ljubljana, Slovenia, April 19-23, 2021*, 2636–2647. ACM / IW3C2.
- Ontanón, S., and Plaza, E. 2010. Amalgams: A formal approach for combining multiple case solutions. In *International Conference on Case-Based Reasoning*, 257–271. Springer.
- Porello, D.; Kutz, O.; Righetti, G.; Troquard, N.; Galliani, P.; and Masolo, C. 2019. A toothful of concepts: Towards a theory of weighted concept combination. In Simkus, M., and Weddell, G. E., eds., *Proc. of the 32nd International Workshop on Description Logics*, volume 2373 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Quine, W. V. O. 1960. *Word & Object*. MIT Press.
- Righetti, G.; Porello, D.; Kutz, O.; Troquard, N.; and Masolo, C. 2019. Pink panthers and toothless tigers: three problems in classification. In Cangelosi, A., and Lieto, A., eds., *Proc. of the 7th International Workshop on Artificial Intelligence and Cognition*, volume 2483 of *CEUR Workshop Proceedings*, 39–53. CEUR-WS.org.
- Righetti, G.; Masolo, C.; Troquard, N.; Kutz, O.; and Porello, D. 2021a. Concept combination in weighted logic. In *Proceedings of the Joint Ontology Workshops 2021, Episode VII*, volume 2969 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Righetti, G.; Porello, D.; Troquard, N.; Kutz, O.; Hedblom, M. M.; and Galliani, P. 2021b. Asymmetric Hybrids: Dialogues for Computational Concept Combination. In *Formal Ontology in Information Systems (FOIS 2021)*, volume 344 of *Frontiers in Artificial Intelligence and Applications*, 81–96. IOS Press. FOIS Best Paper Award.
- Rosch, E. H. 1973. On the internal structure of perceptual and semantic categories. In Moore, T. E., ed., *Cognitive Development and Acquisition of Language*. San Diego: Academic Press. 111–144.
- Ruts, W.; De Deyne, S.; Ameel, E.; Vanpaemel, W.; Verbeemen, T.; and Storms, G. 2004. Dutch norm data for 13 semantic categories and 338 exemplars. *Behavior Research Methods, Instruments, & Computers* 36(3):506–515.
- Schmidt, D.; Pease, A.; Trojahn, C.; and Vieira, R. 2019. Aligning conference ontologies with SUMO: A report on manual alignment via wordnet. In Barton, A.; Seppälä, S.; and Porello, D., eds., *Proc. of the Joint Ontology Workshops 2019, Episode V*, volume 2518 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Silva, V. S.; Freitas, A.; and Handschuh, S. 2016. Word Tagging with Foundational Ontology Classes: Extending the WordNet-DOLCE Mapping to Verbs. In *20th International Conference on Knowledge Engineering and Knowledge Management - Volume 10024, EKAW 2016*, 593–605. Berlin, Heidelberg: Springer-Verlag.
- Speer, R.; Chin, J.; and Havasi, C. 2017. Conceptnet 5.5: An open multilingual graph of general knowledge. In Singh, S. P., and Markovitch, S., eds., *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA*, 4444–4451. AAAI Press.
- Storms, G.; De Boeck, P.; and Ruts, W. 2000. Prototype and exemplar-based information in natural language categories. *Journal of Memory and Language* 42(1):51–73.
- Turvey, M. T. 1992. Affordances and prospective control: An outline of the ontology. *Ecological psychology* 4(3):173–187.
- Veale, T. 2019. From Conceptual Mash-ups to Badass Blends: A Robust Computational Model of Conceptual Blending. In *Computational Creativity*. Springer. 71–89.
- Vinson, D. P., and Vigliocco, G. 2008. Semantic feature production norms for a large set of objects and events. *Behavior Research Methods* 40(1):183–190.
- Völker, J.; Hitzler, P.; and Cimiano, P. 2007. Acquisition of OWL DL axioms from lexical resources. In Franconi, E.; Kifer, M.; and May, W., eds., *In proc. ESWC 2007*, volume 4519 of *Lecture Notes in Computer Science*, 670–685. Springer.
- Yeh, W., and Barsalou, L. W. 2006. The situated nature of concepts. *The American Journal of Psychology* 119(3):349–384.