

Spatial reference in the Semantic Web and in Robotics

Edited by

Aldo Gangemi¹, Verena V. Hafner², Werner Kuhn³,
Simon Scheider⁴, and Luc Steels⁵

1 CNR-ISTC National Research Council – Rome, IT, aldo.gangemi@cnr.it

2 Humboldt-Universität zu Berlin, DE, hafner@informatik.hu-berlin.de

3 University of California – Santa Barbara, US, kuhn@spatial.ucsb.edu

4 Westfälische Wilhelms-Universität Münster, DE,
simon.scheider@uni-muenster.de

5 Free University of Brussels, BE

Abstract

The seminar on “Spatial reference in the Semantic Web and in Robotics” was held from March 30 until April 4, 2014. Seminar participants presented their work related to spatial reference from the viewpoint of Robotics, Spatial Cognition, Geospatial information and the Semantic Web. Groups worked on concrete questions and challenges which were developed during the seminar, some of which resulted in follow up work. This report summarizes the outcomes of the seminar discussions and presents the abstracts of participant talks.

Seminar March 30 to April 4, 2014 – <http://www.dagstuhl.de/14142>

1998 ACM Subject Classification I.2.1 Applications and Expert Systems, I.2.4 Knowledge Representation Formalisms and Methods, I.2.9 Robotics

Keywords and phrases Spatial reference systems, Semantic Web, Robotics, Embodiment, Spatial cognition

Digital Object Identifier 10.4230/DagRep.4.3.181

Edited in cooperation with Saša Bodiroža

1 Executive Summary

Aldo Gangemi

Verena V. Hafner

Werner Kuhn

Simon Scheider

Luc Steels

License © Creative Commons BY 3.0 Unported license

© Aldo Gangemi, Verena V. Hafner, Werner Kuhn, Simon Scheider, and Luc Steels

Motivation

Places (“downtown”), spatial objects (“highway 1”) and localized events (“hurricane Katrina”), are commonly referred to in the Semantic Web. They serve to search for and link to information across domains. Spatial reference systems, such as WGS84¹, allow for representing

¹ The world geodetic system https://en.wikipedia.org/wiki/World_Geodetic_System, a standard reference ellipsoid and coordinate frame for encoding locations on the earth surface.



Except where otherwise noted, content of this report is licensed under a Creative Commons BY 3.0 Unported license

Spatial reference in the Semantic Web and in Robotics, *Dagstuhl Reports*, Vol. 4, Issue 3, pp. 181–201
Editors: Aldo Gangemi, Verena V. Hafner, Werner Kuhn, Simon Scheider, and Luc Steels



Dagstuhl Reports

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

such references as points or regions. This makes them amenable not only for mapping, but also for powerful location-based querying, navigation support and computing.

Spatial references are also fundamental in embodied cognition and robotics. Egocentric and allocentric spatial reference frames underlie robot learning and interaction. Decades of research in cognitive robotics highlight the role of social interaction, joint attention, language games, and visual discrimination games in establishing referents for symbols. The most well-known experiment is that of the Talking Heads. Spatial relations, such as right, front, left, behind, serve to name and identify other objects in a self-organizing vocabulary. Affordance-based cognition is a source of spatial reference in robots as well as in humans. However, so far, this research is only loosely connected to information science and the Semantic Web.

Existing options to localize information in the Semantic Web and in Robotics through coordinate systems cover only limited cases of spatial reference. Humans localize referents in space in many ways, based on different tasks and spatial competencies. For example, the location of a workplace may be linked to people, tasks, and infrastructures. It can be specified in terms of a coordinate system or, alternatively, in terms of containment, connectedness and accessibility in a building; yet another option is to specify it by the possibility to perform certain activities, such as sitting or reading and writing at the workplace.

The seminar

This Dagstuhl Seminar brought together leading international researchers from the Semantic Web, Spatial Cognition, Geo-informatics and Cognitive Robotics to work on the application, synthesis, formal construction, extension, and use of spatial reference systems, identifying challenges and research opportunities. The seminar gathered 27 researchers, 9 from Spatial Cognition and reasoning, 6 from Geo-informatics, 7 from Cognitive Robotics, and 5 from the Semantic Web.

Seminar participants identified a number of concrete links between these communities that are being exploited for future research and development. For example, spatial reference systems of robots and corresponding cognitive spatial concepts can be used in order to describe resources accessible in the world, and Semantic Web technology to publish those descriptions for information access. Locations can be described in ways which are more closely related to humans, based on qualitative relations and environmental referents, and for environments which are difficult to localize by a GPS. In this way, it becomes possible to share location descriptions among humans and robots and thus to localize resources of interest (e. g. rooms, people, places) published in the Web of data. Vice versa, spatial referents and descriptions in the Semantic Web may guide robots towards accessible things in the world. Robots may function as embodied surrogates of human observers exchanging information on the Web of Data encoded in terms of their own reference systems.

2 Table of Contents

Executive Summary

Aldo Gangemi, Verena V. Hafner, Werner Kuhn, Simon Scheider, and Luc Steels . 181

Overview of Talks


Spatial Descriptions in Formal and Natural Languages <i>Brandon Bennett</i>	185
Visuo-Spatial Reasoning for Computational Cognitive Systems <i>Mehul Bhatt</i>	185
Gestures in Human-Robot Interaction <i>Saša Bodiřoža</i>	187
Why Spatial Reference should eschew projections <i>Nicholas Chrisman</i>	188
“The kids are in the kitchen!!” Artificial spaces as place-reference systems <i>Helen Couclelis</i>	188
An Analysis of the Semantics of Spatial Representation and Relations <i>Anne Cregan</i>	189
A preliminary sketch of some definitions for “Data”, “Information”, “Knowledge” and “Wisdom” <i>Anne Cregan</i>	189
Spatial Frames across Language and the Web <i>Aldo Gangemi</i>	190
Cross-World Identity in Conceptual Spaces <i>Giancarlo Guizzardi</i>	191
Attention Models in Robotics <i>Verena V. Hafner</i>	191
Ontology Design Patterns for Ocean Science Data Discovery <i>Pascal Hitzler</i>	192
Spatial Referencing <i>Werner Kuhn</i>	192
Spatial Cognition in Robotics <i>Bruno Lara Guzman</i>	193
Observation of Human Activity in Intelligent Space – Human-Object/Environment Interaction <i>Mihoko Niitsuma</i>	194
Representing structural, functional and organizational dimensions of indoor spaces <i>Kai-Florian Richter</i>	195
Decontextualizing spatial words in the process of language acquisition <i>Katharina J. Rohlfing</i>	195
Grounding Spatial Language in Sensorimotor Experience <i>Yulia Sandamirskaya</i>	195

Treasure maps. Spatial reference systems for the common man <i>Simon Scheider</i>	196
Place-based GIS <i>Stephan Winter</i>	196
Qualitative Relations vs. Locative Expressions <i>Diedrich Wolter</i>	197
Working Groups	
Spatial reference design pattern <i>Aldo Gangemi, Pascal Hitzler, and Simon Scheider</i>	197
Draft OWL specification of Spatial Reference Pattern <i>Aldo Gangemi, Pascal Hitzler, and Simon Scheider</i>	198
InnoCentive Challenge Submission on Gazetteers for Place Descriptions <i>Werner Kuhn</i>	198
Panel Discussions	
The discussions	198
The outcomes	199
Participants	201

3 Overview of Talks

3.1 Spatial Descriptions in Formal and Natural Languages

Brandon Bennett (*University of Leeds, GB*)

License  Creative Commons BY 3.0 Unported license
© Brandon Bennett

This paper explores the complex mapping between language and meaning and the repercussions this has for formalising natural language semantics. The general problem will be investigated by consideration of the domain of spatial descriptions, especially those that describe a spatial relationship between two entities. I argue that the meanings of spatial expressions of natural language can be modelled in terms of a cluster of distinct formal definitions, each of which corresponds to an artificially precise version of the informal natural concept. Moreover, I show how one can derive a probability distribution over possible formal interpretations of natural language vocabulary and phrases by analysing their occurrence in a corpus that is representative of typical language use.

3.2 Visuo-Spatial Reasoning for Computational Cognitive Systems

Mehul Bhatt (*Universität Bremen, DE*)

License  Creative Commons BY 3.0 Unported license
© Mehul Bhatt

Joint work of Bhatt, Mehul; Schultz, Carl

We pursue *visuo-spatial representation and reasoning* from the viewpoint of the research areas of *artificial intelligence*, *commonsense reasoning*, and *spatial cognition and computation*.

We propose **declarative spatial reasoning** as the ability to (declaratively) specify and solve real-world problems related to geometric and qualitative visuo-spatial representation and reasoning [2]. The problems that we address in this context encompass both specialist and everyday instances identifiable in a range of *cognitive technologies* and *spatial assistance systems* where spatio-linguistic conceptualisation & background knowledge focussed visuo-spatial cognition and computation are central [3].

As a first step toward the systematic development of the declarative spatial reasoning method, we have initiated formalisations of *space* and *spatial reasoning* within constraint logic programming [2, 6, 7]. We have developed **CLP(QS)**, a *declarative spatial reasoning system* capable of modelling and reasoning about qualitative spatial relations pertaining to multiple spatial domains, i. e., one or more aspects of space such as topology, and intrinsic and extrinsic orientation, size, distance etc. With CLP(QS), users and application developers may freely mix object domains (i. e., points, line-segments, and regions) with the available spatial domains. CLP(QS) also offers mixed geometric-qualitative spatial reasoning capabilities, and in its current form, basic quantification support offering the means to go back from qualitative relations to the domain of precise quantitative information.

The emphasis in CLP(QS) is on the seamless integration of declarative visuo-spatial (computational) problem-solving capabilities within large-scale hybrid AI systems, and cognitive (interaction) technologies. Currently, integration is achieved via the medium of *logic programming* – specifications in the form of (domain) *facts* and *rules* consisting of mix of, for instance, background semantic or conceptual knowledge, spatio-temporal knowledge, and knowledge about action and dynamics. The general concept of declarative spatial reasoning

lends itself to re-interpretations and extensions with other perspectives such as diagrammatic representations.

CLP(QS) marks a clear departure from other (relational-algebraically based) spatial reasoning methods / tools by its use of the constraint logic programming framework for formalising the semantics of mixed geometric and qualitative spatial representation and reasoning. The approach has demonstrated applicability in several domains, most recent examples being *architectural design cognition* [5], *cognitive vision* [4], *geospatial information systems* [1].

The CLP(QS) system is also being designed and used as a pedagogical tool to be used as part of university based courses at the interface of Artificial Intelligence, Knowledge Representation and Reasoning, Cognitive Systems, and Spatial Informatics.

CLP(QS): A Declarative Spatial Reasoning System. <http://www.spatial-reasoning.com>

References

- 1 M. Bhatt and J. O. Wallgruen. Geospatial narratives and their spatio-temporal dynamics: Commonsense reasoning for high-level analyses in geographic information systems. *ISPRS International Journal of Geo-Information.*, Special Issue on: *Geospatial Monitoring and Modelling of Environmental Change*, *ISPRS International Journal of Geo-Information*, 3 (1):166–205, 2014. ISSN 2220-9964. 10.3390/ijgi3010166. URL <http://www.mdpi.com/2220-9964/3/1/166>.
- 2 M. Bhatt, J.H. Lee, and C. Schultz. CLP(QS): A Declarative Spatial Reasoning Framework. In *Proceedings of the 10th international conference on Spatial information theory, COSIT'11*, pages 210–230, Berlin, Heidelberg, 2011. Springer-Verlag. ISBN 978-3-642-23195-7.
- 3 M. Bhatt, C. Schultz, and C. Freksa. The ‘Space’ in Spatial Assistance Systems: Conception, Formalisation and Computation. In T. Tenbrink, J. Wiener, and C. Claramunt, editors, *Representing space in cognition: Interrelations of behavior, language, and formal models. Series: Explorations in Language and Space*. 978-0-19-967991-1, Oxford University Press, 2013a.
- 4 M. Bhatt, J. Suchan, and C. Schultz. Cognitive Interpretation of Everyday Activities – Toward Perceptual Narrative Based Visuo-Spatial Scene Interpretation. In M. Finlayson, B. Fisseni, B. Loewe, and J.C. Meister, editors, *Computational Models of Narrative (CMN) 2013., a satellite workshop of CogSci 2013: The 35th meeting of the Cognitive Science Society*, Dagstuhl, Germany, 2013b. OpenAccess Series in Informatics (OASiCs).
- 5 M. Bhatt, C. Schultz, and M. Thosar. Computing narratives of cognitive user experience for building design analysis: Kr for industry scale computer-aided architecture design. In T. Eiter, C. Baral, and G.D. Giacomo, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the 14th International Conference, KR 2014*, 2014. <http://kr.org/KR2014/>. To appear.
- 6 C. Schultz and M. Bhatt. Towards a Declarative Spatial Reasoning System. In *ECAI 2012 – 20th European Conference on Artificial Intelligence, Montpellier, France*, 2012.
- 7 C. Schultz and M. Bhatt. Declarative Spatial Reasoning with Boolean Combinations of Axis-Aligned Rectangular Polytopes. In *ECAI 2014 – 21st European Conference on Artificial Intelligence, Prague*, 2014. to appear.

3.3 Gestures in Human-Robot Interaction

Saša Bodiřoža (*Humboldt-Universität zu Berlin, DE*)

License © Creative Commons BY 3.0 Unported license
© Saša Bodiřoža

Joint work of Bodiřoža, Saša; Hafner, Verena V.; Lara, Bruno; Doisy, Guillaume; Jevtić, Aleksandar

Intuitive human-robot interaction relies on multimodal communication between humans and robots by means of symbols which are easily understood by participants. This work focuses on the use of gestures in human-robot interaction and consists of two steps: development of gesture vocabularies and gesture recognition.

Gesture vocabularies represent mappings between actions and gestures. Usually, they are constructed by system designers and represent 1–1 mappings. This is far from the case in the real world, where these are n–n mappings. Therefore, they need to be defined through observation of the associations people make between particular actions and gestures. To construct a human gesture vocabulary, a user survey is performed to collect gestures which people usually associate with pre-selected actions [1]. The results present an important input before the design stage of the gesture recognizer. A robot gesture vocabulary is developed in a similar manner, where participants are asked to rank videos of pre-recorded robot gestures, which are based on the gestures that humans perform. Resulting gestures might not be the best for a particular person. In order to make them more general, interactive evolution of these gestures is performed to come up with better variant of these gestures. Every time a human being gesture, they tend to to perform the same gesture with slight variances in direction, size, velocity, starting and ending positions of the stroke, location where they gesture and the orientation between the person who gestures and the observer.

One way to make gesture recognition invariant to these variables is to pre-process trajectories, before the recognition, so that these features are removed. A gesture recognition algorithm, based on dynamic time warping, presents a way to recognize gestures which are performed with varying velocity [2]. Inputs are gesture trajectories, which are obtained using a RGB-D camera and transformed from the robot's frame of reference to that of the person. Furthermore, the trajectories are normalized and aligned with the trained gestures to make the recognizer robust to gestural size and location. Finally, the gesture recognition is made more robust to by including an invariance to the direction of the gesture. Both algorithms are trained using one sample per gestural class. This leads to unrestricted gesture recognition, enabling its application in real world scenarios. However, the above mentioned features (e. g. direction and size) are relevant for the understanding of some gestures and a disambiguation framework is presented, which is trained to disambiguate gestures based on the output of the recognizer and these features where they are relevant.

Work on motion control learning will be presented through two experiments. Both present action execution systems that rely on learned sensorimotor schemes. These schemes are learned as a product of the interaction of an agent with its environment. In the first experiment, a mobile agent learns an association between changes in its sensory perception and the random movements it performs (so called motor babbling) [3]. Once it has this acquired knowledge, the agent is then capable of performing a mirror action to match an observed gesture. This is seen as a first step toward learning motor control strategies for a robot control task. In the second experiment, it learns associations between changes in its sensory perception and its movement, guided by the demonstrator. After learning, it is capable of executing the necessary motor commands to go to a location where the demonstrator is pointing.

References

- 1 Saša Bodiroža, Helman I. Stern, and Yael Edan. *Dynamic gesture vocabulary design for intuitive human-robot dialog*. In *Proceedings of the 7th ACM/IEEE international conference on Human-robot interaction, HRI'12*, pages 111–112, New York, NY, USA, 2012. ACM.
- 2 Saša Bodiroža, Guillaume Doisy, and Verena Vanessa Hafner. *Position-invariant, real-time gesture recognition based on dynamic time warping*. In *Proceedings of the 8th ACM/IEEE international conference on Human-robot interaction, HRI'13*, pages 87–88, Piscataway, NJ, USA, 2013. IEEE Press.
- 3 Saša Bodiroža, Aleksandar Jevtić, Bruno Lara, and Verena V. Hafner. *Learning the relation of motion control and gestures through self-exploration*. In *Proceedings of Robotics Challenges and Vision Workshop, at Robotics: Science and Systems*, Berlin, Germany, 2013.

3.4 Why Spatial Reference should eschew projections

Nicholas Chrisman (RMIT University – Melbourne, AU)

License  Creative Commons BY 3.0 Unported license
© Nicholas Chrisman

A large proportion of GIS databases use a local reference system based on linear measurements (metres) on a projection plane of some form. This practice creates systematic errors which are not corrected in routine practice.

Projection error is nothing new, but it is no longer required as computation has become cheaper. It is time to stop this practice and move to the round (ellipsoidal) Earth as a reference system. Current dynamic reference frameworks (such as ITRF) provide a more solid and sustainable basis for spatial reference.

3.5 “The kids are in the kitchen!!” Artificial spaces as place-reference systems

Helen Couclelis (University of California – Santa Barbara, US)

License  Creative Commons BY 3.0 Unported license
© Helen Couclelis


Joint work of Couclelis, Helen; Howarth, Jeffrey

Emergency personnel entering a house they have never seen before will turn towards the living room or dining area rather than the bedrooms in order to reach the kitchen and rescue the children. This is a simple inference to make because a house, like all artificial things, is made for a purpose: here, the purpose of supporting a household’s everyday activities, and more specifically in the case of the kitchen, the activity of storing and preparing food, which is closely linked to the activity of eating (but not of sleeping). The same general principle applies to artificial or artificially configured spaces at any scale, from that of the layout of the kitchen itself to those of farms or harbors or national parks, in that each of these spaces must function so as to enable the activities it is meant to support. It turns out that the logic of the mapping between human activities and artificial spatial structures can be modeled with a degree of accuracy sufficient for many applications, allowing inferences of structure from activity and conversely. In geographic-scale artificial spaces that same logic connecting functional relations and spatial structure conveys properties of place-reference

systems. Indeed, by designating a place with a name that reflects its purpose (the kitchen, the auditorium, the barn, and so on), one provides valuable information on the likely location of other places functionally related to the former.

3.6 An Analysis of the Semantics of Spatial Representation and Relations

Anne Cregan (Intersect – Sydney, AU)

License  Creative Commons BY 3.0 Unported license
© Anne Cregan


I'd like to present my analysis of the semantics of spatial representation and relations through a number of examples. These will highlight the importance of asking these questions and being able to formulate good answers for them:

1. What is the frame of reference, i. e., the space to be represented? What are its underlying dimensions and properties?
2. What is our motivation for representing this space, i. e., what are we actually trying to do? What kind of representation is most appropriate for this?
3. How do we intend to “ground” the representation to the space, i.e. how should we establish correspondence points, lines or other features which connect the representation to the space being represented?
4. When is a space really a network?
5. What is the difference between an absolute and relative reference to a location within a space?
6. What spatial relations are meaningful in the space and how should we represent them?
7. Factors pertaining to agents and objects operating in the space
8. Factors pertaining to solving problems of logistics, optimisation etc. as they relate to the space

As this analysis works through what I believe are some of the most fundamental issues and lays out a framework for further discussions, it would best be placed early in the schedule.

3.7 A preliminary sketch of some definitions for “Data”, “Information”, “Knowledge” and “Wisdom”

Anne Cregan (Intersect – Sydney, AU)

License  Creative Commons BY 3.0 Unported license
© Anne Cregan

Stimulated by discussion during the seminar, this is a preliminary sketch of how one might define these terms.

Some of the questions raised are:

- Are these terms representing the categories that are discrete or do they fall on a spectrum?
- If they are values falling on a single dimension, then what is the underlying dimension?
- Is wisdom on an orthogonal dimension to the data-to-knowledge spectrum?

3.8 Spatial Frames across Language and the Web

Aldo Gangemi (CNR – Rome, IT)

License  Creative Commons BY 3.0 Unported license
© Aldo Gangemi

URL <http://stlab.istc.cnr.it/documents/slides/kp-spatialref-lod.pdf>

My presentation is an introduction to frames and knowledge patterns, as they can be applied – jointly with semantic web, linked data, and knowledge extraction techniques – to empirical investigations of Spatial Knowledge Representation (SKR).

Frames are task-oriented invariants in conceptualizing the world. They are related to both cognition (conceptual schemas) and problem solving (design patterns). Frames appear in data, ontologies, conceptual schemas, language, web and interaction formats, etc. A structure-neutral, broad notion of frames is “Knowledge Pattern” (KP), a notion introduced firstly by Robert de Beaugrande in the seventies to generalize over frames, scripts, schemas, as they were used in AI, linguistics, and cognitive science. Knowledge patterns have been formalized firstly by Peter Clark in late nineties, and applied to the semantic web by Aldo Gangemi and Valentina Presutti. In this late notion, knowledge patterns are meant for reusability and good design practices: they are reusable successful solutions to recurrent modeling problems. A useful feature of KPs is that they are invariant across several representation layers (e.g. natural language-data-schema; schema-data, etc.): therefore KPs are supposed to enable interoperability and knowledge discovery across arbitrary representation languages and formats. Knowledge patterns emerge as top-down collections of good practices, or as empirical findings from data, text, ontologies, etc.

SKR is core to any linked data design project. Currently all important LOD bubbles contain spatial knowledge: DBpedia (about 800 thousand spatial entities and about 3 million spatial facts), Yago (about 900 thousand spatial entities), Geonames, GeoLinkedData. Schemas have been designed by reusing existing SKR, notably the GeoSPARQL ontology, which incorporates part of the RCC8 spatial algebra. On one hand, linked spatial data are however incomplete and prone to inconsistencies, e.g. 3646 Yago spatial entities are typed so that inconsistencies emerge as soon as plausible disjointness axioms are used to reason over that ontology. On the other hand, the Web is an ideal platform for experimenting with SKR. Besides the amount of public data and their public, reusable identities, the Web manifests original and interesting SKR problems: user-centric space, web-specific space, spatial predicates embedding events or other relations.

When natural language is put into the picture, things become even more exciting: Levinson’s Spatial Reference Systems (SRS: Absolute, Intrinsic, Relative) map in a complex way against linguistic constructions: locally constructed SRS, metonymic SRS, metaphoric SRS. Levinson’s SRS and SRS emerging from the Web and from particular linguistic constructions correspond to novel Spatial Knowledge Patterns (SKP). Some of them can be found in top-down KP repositories such as FrameNet, others can be discovered from data and text. Others are found or emerge out of text. Deep machine reading such as the one performed by the FRED tool (<http://wit.istc.cnr.it/stlab-tools/fred>) allows to combine knowledge extraction based on NLP algorithms with existing linked data, so enabling a human-like contextual reading of text including spatial references.

In order to apply SKP to SRS, we need powerful KPs that abstract from the high variability of SRS phenomena. A good candidate for reuse is the Descriptions and Situations (D&S) pattern, which is able to pair “duper” relations, typically a factual one (e.g. a situation with entities at places during events with characteristics), and a descriptive one (e.g. a

description of a relative SRS used to understand that situation). Applying D&S to SKR on the Web requires reification of n-ary relations, “punning” of predicates, and dynamic typing of named graphs for RDF serialization. The SRS pattern drafted during the seminar is a case of SKP, and is partly inspired by D&S.

Please refer to the slides of my presentation at the seminar for details.

3.9 Cross-World Identity in Conceptual Spaces

Giancarlo Guizzardi (UFES – Vitoria, BR)

License © Creative Commons BY 3.0 Unported license
© Giancarlo Guizzardi

Main reference G. Guizzardi, “Logical, Ontological and Cognitive Aspects of Objects Types and Cross-World Identity with applications to the theory of Conceptual Spaces,” to appear in P. Gardenfors, F. Zenker, eds., “Applications of Geometric Knowledge Representation,” Synthese Library, Dordrecht, Springer; forthcoming.

The theory of Conceptual Spaces put forth by the Swedish philosopher and cognitive scientist Peter Gärdenfors proposes a geometrical model for representing and reasoning with Categories and Properties. The theory has been shown to serve as a valuable tool in philosophy, cognitive science, linguistics and computer science, in a number of applications ranging from foundations for conceptual modeling to Robotics. However, recently, the theory has suffered criticisms due to its insufficiency in supporting the structure of judgments and, in particular, its unsatisfactory treatment of cross-world identity and persistence for enduring individuals. In this work, I briefly discuss a philosophically and cognitively well-founded theory of object categories as well as a system of Sortal Intensional Logic (a Modal Logic with Individual Concepts and Sortal-restricted quantification) derived from this theory. Moreover, I illustrate how these results can be used to address the aforementioned limitations of the theory of Conceptual Spaces.

3.10 Attention Models in Robotics

Verena V. Hafner (Humboldt-Universität zu Berlin, DE)

License © Creative Commons BY 3.0 Unported license
© Verena V. Hafner

Joint work of Hafner, Verena V.; Schillaci, Guido; Bodiroža, Saša

Main reference G. Schillaci, S. Bodiroža, V. V. Hafner, “Evaluating the Effect of Saliency Detection and Attention Manipulation in Human-Robot Interaction,” *International Journal of Social Robotics*, 5(1):139–152, 2013; available open access.

URL <http://dx.doi.org/10.1007/s12369-012-0174-7>

The ability to share the attention with another individual is essential for having intuitive interaction and to communicate about spatial events. Two relatively simple, but important prerequisites for this, saliency detection and attention manipulation by the robot, are identified. By creating a saliency based attentional model combined with a robot ego-sphere [2] and by adopting attention manipulation skills, the robot can engage in an interaction with a human and start an interaction game including objects as a first step towards a joint attention [1].

References

- 1 Kaplan, F. and Hafner, V.V. (2006), The Challenges of Joint Attention, *Interaction Studies*, 7:2, pp. 135–169
- 2 Schillaci, G., Bodiřoža, S. and Hafner, V.V. (2013), Evaluating the Effect of Saliency Detection and Attention Manipulation in Human-Robot Interaction, *International Journal of Social Robotics*, Springer, Volume 5, Issue 1 (2013), pages 139–152, DOI: 10.1007/s12369-012-0174-7 OPEN ACCESS.

3.11 Ontology Design Patterns for Ocean Science Data Discovery

Pascal Hitzler (Wright State University – Dayton, US)

License © Creative Commons BY 3.0 Unported license
© Pascal Hitzler

Joint work of Hitzler, Pascal; Krisnadhi, Adila; Arko, Robert; Carbotte, Suzanne; Chandler, Cynthia; Cheatham, Michelle; Finin, Timothy; Janowicz, Krzysztof; Narock, Thomas; Raymond, Lisa; Shepherd, Adam; Wiebe, Peter

EarthCube is a major effort of the National Science Foundation to establish a next-generation knowledge architecture for the broader geosciences. Data storage, retrieval, access, and reuse are central parts of this new effort. Currently, EarthCube is organized around several building blocks and research coordination networks. The NSF EarthCube OceanLink project is currently under way to integrate the two major U.S. ocean science repositories, BCO-DMO and R2R, using a flexible approach based on ontology design patterns, which is set to scale to significant breadth and depth. Ontology design patterns are the method of choice for this integration.

References

- 1 Adila Krisnadhi, Robert Arko, Suzanne Carbotte, Cynthia Chandler, Michelle Cheatham, Timothy Finin, Pascal Hitzler, Krzysztof Janowicz, Thomas Narock, Lisa Raymond, Adam Shepherd, Peter Wiebe. *An Ontology Pattern for Oceanographic Cruises: Towards an Oceanographer's Dream of Integrated Knowledge Discovery*. OceanLink Technical Report 2014.1. Available from <http://www.oceanlink.org/>

3.12 Spatial Referencing

Werner Kuhn (University of California – Santa Barbara, US)

License © Creative Commons BY 3.0 Unported license
© Werner Kuhn

I gave a very brief introduction to the basic ideas underlying spatial referencing on the earth. My framing of the topic suggested ways to generalize from geographic to other spaces and from space (and time) to measurements in general.

3.13 Spatial Cognition in Robotics

Bruno Lara Guzman (Universidad Autónoma del Estado de Morelos, MX)

License © Creative Commons BY 3.0 Unported license
© Bruno Lara Guzman

Intelligence and thinking are mental phenomena essentially linked to both our bodies and the environmental conditions in which we are immersed.

We understand the physical world according to our own experience. The faculties, capabilities and skills to dynamically interact with the world, which as adult humans we possess, emerge through a long process of tuning and rehearsing of sensori-motor schemes.

A very important example is the acquisition of the sensori-motor schemes that code for the capabilities and reaches of our body. This set of schemes provide us with many cognitive tools, among them the knowledge and coding of our body map, essential for among other tasks navigating around the environment.

Throughout its short history, artificial intelligence research has witnessed at least one major paradigm shift, namely from a cognitivist perspective, which conceived cognition as a amodal symbol crunching activity to a new perspective emphasizing the role of the body and environmental structures in cognition.

In the last decades, together with the rest of the sciences studying cognition the shift in paradigm has been towards the rediscovery of the importance the body of agents has on the development of cognition.

In what has become known as embodied cognition, it is widely accepted now that, to be able to understand and replicate intelligence, it is necessary to study agents in their relation to their environment [1]. The subjects of study should be agents with a body, that learn through interaction with their environment and that this learning should be a developmental process.

Pursuing the same aims, cognitive robotics takes its inspiration from studies of cognitive development in humans. Artificial agents or robots, by having a body and being situated in the real world, assure a continuous and real-time coupling of body, control and environment.

Theories of grounded cognition have also had a great impact on this quest. In general, these theories reject the use of modal symbols for the representation of knowledge, focusing on the role of the body for its acquisition [5]. More importantly for our purposes, grounded cognition focuses on the role of internal simulations of the sensorimotor interaction of agents with their environment.

For the advocates of grounded cognition modal simulations, such as recreations of perceptual, motor and introspective states, are important components that allow the development of high cognitive abilities. Such recreations could account for the off-line characteristics of cognition, in which internal simulations of sensory-motor cycles are executed.

Studies that account for these theories provide experimental results that support the importance of internal models and multimodal representations to form the ground capable of supporting the whole cognition scaffolding.

In the quest of a basic internal simulation mechanism *forward and inverse models* have been proposed [3]. A forward model is an internal model which incorporates knowledge about sensory changes produced by self-generated actions of an agent. Given a sensory situation S_t and a motor command M_t (intended or actual action) the forward model predicts the next sensory situation S_{t+1} . While forward models (or predictors) present the causal relation between actions and their consequences, inverse models (or controllers) perform the opposite transformation providing a system with the necessary motor command (M_t) to go from a current sensory situation (S_t) to a desired one (S_{t+1}).

Forward and inverse models become central players in cognition, as they naturally fuse together different sensory modalities as well as motor information providing agents with multimodal representations [2].

We believe that the joint and coordinate action of both forward and inverse models gives an agent a practical sense of situations and can even account for subjective experience as a ground for consciousness [4].

In this new perspective internal models have become a keypoint, given the capabilities they allot agents. We believe that these type of models, specially internal and forward models, have not been fully studied and their usefulness not properly exploited. We present a short review of some of the most well-known architectures and implementations in the area in the search to identify and try fill some of the gaps in their study.


In our research we use these models to let agents learn the spatial characteristics of their own bodies, the environment around them and the sensorimotor associations arising from the interaction among these two.

References

- 1 Margaret Wilson, *Six Views of Embodied Cognition*, Psychonomic Bulletin & Review. 2002.
- 2 Wilson, M. and Knoblich, G. *The case for motor involvement in perceiving conspecifics* Psychological bulletin. Volume 131, pp. 460–473. 2005.
- 3 D. M. Wolpert, Z. Ghahramani and M. I. Jordan. *An internal model for sensorimotor integration*. Science, 5232, 269. 1995.
- 4 Kiverstein, J., *Could a Robot Have a Subjective Point of View?* Journal of Consciousness Studies, 2007, 14, 127–139.
- 5 Barsalou, L. W. *Grounded Cognition*. Annual Reviews Psychology, 2008, 59, 617–645.

3.14 Observation of Human Activity in Intelligent Space – Human-Object/Environment Interaction

Mihoko Niitsuma (Chuo University, JP)

License  Creative Commons BY 3.0 Unported license
© Mihoko Niitsuma

Main reference M. Niitsuma, K. Yokoi, H. Hashimoto, “Describing Human-Object Interaction in Intelligent Space,” in Proc. of the 2nd Int’l Conf. on Human System Interaction (HSI’09), pp. 392–396, IEEE, 2009.

URL <http://dl.acm.org/citation.cfm?id=1689359.1689428>

This talk presents an observation system of human-object interaction in the Intelligent Space (iSpace). Object information is necessary to describe events and human activities which happen in environments. Especially, names, colors, size and shapes of objects can be described manually because we can consider that the information will not change. On the other hand, information such as locations of objects, frequency of use of the objects, the users and motion patterns while using the objects can not be described manually because they depend on individuals and contexts when the objects are used. Therefore, we decided to obtain these kinds of object information through observation of human-object interaction.

3.15 Representing structural, functional and organizational dimensions of indoor spaces

Kai-Florian Richter (Universität Zürich, CH)

License © Creative Commons BY 3.0 Unported license
© Kai-Florian Richter

Joint work of Richter, Kai-Florian; Winter, Stephan; Satosa, Sigit

Main reference K.-F. Richter, S. Winter, S. Santosa, “Hierarchical Representations of Indoor Spaces,”
Environment and Planning B: Planning and Design, 38(6):1052–1070, 2011.

URL <http://dx.doi.org/10.1068/b37057>

People’s mental organization of their spatial knowledge is (often) hierarchical. For example, place descriptions have been shown to have hierarchical structure. Such hierarchical structures provide efficient means to if needed refine descriptions (e. g., to disambiguate) or to coarsen them (e. g., if unsure about the details). Also, different (groups of) people have different views (conceptualizations) of a space, for example, a building, that is reflected in differences in their mental representation and the way they talk about that space.

We developed a hierarchical representation of indoor spaces that is based on image schemata and captures structural, functional and organizational dimensions of a space. The representation allows accounting for different use roles (user groups). Still, there is further research required, for example, to automatically convert floor plans or to properly link geometry with the representation.

3.16 Decontextualizing spatial words in the process of language acquisition

Katharina J. Rohlfing (Universität Bielefeld, DE)

License © Creative Commons BY 3.0 Unported license
© Katharina J. Rohlfing

For establishing a representation from (1) a scene on the one hand and (2) a spatial preposition describing this scene on the other hand, different memory processes seem to be responsible. While a first exposure to a new preposition is “fast mapped” in which process contextual cues are crucial, children need to be exposed to different situations in order to “slow map” the new word. In this latter process, the learner becomes able to apply the acquired word in situations that are unfamiliar and provide little cues.

3.17 Grounding Spatial Language in Sensorimotor Experience

Yulia Sandamirskaya (Ruhr-Universität Bochum, DE)

License © Creative Commons BY 3.0 Unported license
© Yulia Sandamirskaya

Dynamic Field Theory (DFT) is a mathematical and conceptual framework, in which emergence of cognitive processes from continuous in time and in space neuronal dynamics may be modelled. Elementary cognitive functions, such as memory formation, feature binding, decision making, and coordinate frame transformations may be expressed in this framework using Dynamic Neural Fields, organised in autonomous neural-dynamic architectures. In

my talk, I demonstrate how the principles, structures, and dynamics of DFT may be put together in an architecture, capable of grounding spatial language in a visually perceived scene, answering queries about the scene, and directing actions at objects in the scene by spatially referencing them.

3.18 Treasure maps. Spatial reference systems for the common man

Simon Scheider (Universität Münster, DE)

License  Creative Commons BY 3.0 Unported license
© Simon Scheider

Spatial reference systems are not only a basis for producing maps, they are also a powerful tool for localizing objects in everyday human communication. The human localization practice, as applied in orientation tasks, is very different from the practice underlying geographic coordinate systems, as used in current IT systems or in GPS. While geographic coordinates are grounded in the earth as referent, human localization practice resembles a “treasure map”, i. e., a description of locations relative to various kinds of referent objects and spatial operations. One way to analyze such systems is to look at localization games, such as “orienteering”. The international orienteering federation proposes a standard syntax for describing the locations of control flags in an orienteering terrain. This standard uses a well defined range of definite descriptors, types and qualities for selecting reference objects (relatums) in the landscape, as well as a choice of relations to determine locations relative to objects, based on cardinal and vertical directions, imposed orderings and object morphology. The benefit of such reference systems is that they work indoors as well as outdoors, at least in principle, and that they reflect human localization practice. In future work, one could investigate how to formalize orienteering systems, and how to support spatial computations as well as translations to geographic coordinate reference systems. A concrete challenge is to build a digital orienteering system, which helps players compare and retrieve control descriptions, and users describe locations precisely, in a manner which comes close to their habit. From a larger perspective, such a system could demonstrate how to make available spatial reference systems for the common man in the Semantic Web, in order to describe locations of arbitrary resources.

3.19 Place-based GIS

Stephan Winter (The University of Melbourne, AU)

License  Creative Commons BY 3.0 Unported license
© Stephan Winter

People reason about space by places, and configurations of places. Configurations of known (or more salient) places form the spatial reference system in language: An expression “the coffee shop opposite the library” assumes the library (relatum) is shared spatial knowledge, and positions the unknown or ambiguous coffee shop (locatum) to this relatum.

Language typically remains qualitative, i. e., collecting corpora of natural language descriptions would provide arbitrary large sets of triples <locatum, reference, relatum> (for binary spatial relationships). These triplets do not only provide a relationship, they also distinguish a feature that is positioned from a feature that allows positioning. The set

of relata must be the features commonly known, or at least shared between speaker and recipient in the individual communication situations. In this sense, the set of relata seems to determine a spatial reference frame.

This talk poses the challenge to build place-based GIS: GIS where the frame of reference comes from configurations of known places, instead of a metric externally defined spatial reference system.

3.20 Qualitative Relations vs. Locative Expressions

Diedrich Wolter (Universität Bamberg, DE)

License © Creative Commons BY 3.0 Unported license

© Diedrich Wolter

Joint work of Mast, Vivien; Wolter, Diedrich

Main reference V. Mast, D. Wolter, “Context and Vagueness in REG,” in Proc. of the 2013 Workshop on Production of Referring: Bridging the gap between cognitive and computational approaches to reference (PRE-CogSci’13), 6 pages, 2013.

URL <http://pre2013.uvt.nl/pdf/mast-wolter.pdf>

Qualitative spatial reasoning is a subfield of AI research involved with symbolic representations that aims to capture common sense spatial concepts. Representations in qualitative spatial reasoning are based on finite sets of relations that categorize spatial properties by comparison. Among other applications, qualitative spatial representations are widely acknowledged to provide a basis for natural language semantics as well relations need to exhibit specific features to make reasoning efficient. However, we recently observed properties of locative expressions that do not align with qualitative representations studied so far. This talk will contrast some of the computational requirements with language use.

4 Working Groups

4.1 Spatial reference design pattern

Aldo Gangemi, Pascal Hitzler, and Simon Scheider


License © Creative Commons BY 3.0 Unported license

© Aldo Gangemi, Pascal Hitzler, and Simon Scheider

During the seminar, one group worked on an ontology design pattern which captures central aspects of Stephen Levinson’s typology of reference frames (relative, absolute and intrinsic) and makes them formally explicit, so that one can distinguish the different ways how humans and robots refer to space. The idea is to design an ontology pattern which can be used to encode the different ways how agents actually construct spatial referents. Examples for the latter are: walk for 30 minutes and you’re there; in front of the house; outside the house; north of the oak tree in the park; this theater, room 2, seat 4A; the parking lot next to the cinema; between the tree and the well; at the foot of the cliff; the tree uphill from the house. In robotics, spatial referents of this kind play a major role not only in sensory-motor coordination, but also in robot interaction and learning. Can we generate a general design pattern that allows to document how agents construct spatial referents? This would enable automatic descriptions of locations visited by a robot or a human, which could be published and compared in the Semantic Web as a communication layer. The group plans to write a paper about this, in which concrete spatial reference frames used in robotic experiments are taken as empirical data on which the pattern is tested.

4.2 Draft OWL specification of Spatial Reference Pattern


Aldo Gangemi, Pascal Hitzler, and Simon Scheider

License  Creative Commons BY 3.0 Unported license
 © Aldo Gangemi, Pascal Hitzler, and Simon Scheider
URL <http://www.ontologydesignpatterns.org/cp/owl/spatial-reference.owl>

This is a draft of the OWL spec for the Spatial Reference Pattern, as designed in Dagstuhl by Aldo Gangemi, Pascal Hitzler, and Simon Scheider. It is still very incomplete in terms of axioms, domain coverage, and annotations. No links are yet provided to existing vocabularies.

4.3 InnoCentive Challenge Submission on Gazetteers for Place Descriptions

Werner Kuhn (University of California – Santa Barbara, US)

License  Creative Commons BY 3.0 Unported license
 © Werner Kuhn
Joint work of Kuhn, Werner; Ballatore, Andrea; Janowicz, Krzysztof

We propose an annual computational cognition competition to build gazetteer services that interpret geographic place descriptions. Examples for such descriptions are “the hill along Grattan street” or “in front of the door of the smithy in Port Khazard”. Locating unnamed places is increasingly important and helps to interpret and integrate big data, in applications ranging from the digital humanities through physical infrastructure maintenance to security. Solutions shall map described places from a corpus as points, lines, or polygons at an adequate scale, using an online map service; they shall also return the geometries encoded as well-known text , at a resolution corresponding to their estimated precision.

5 Panel Discussions

5.1 The discussions

A number of concrete questions evolved during the seminar and were investigated in group work. During the afternoon sessions, participants split into groups discussing the following topics:

- How to construct a digital orienteering system? Such a system may help people who plan orienteering games² generate and share orienteering descriptions for a given terrain, generate automatic cartographic mappings from orienteering descriptions, or it may help robots play this game.
- What are criteria for selecting spatial referents? There seems to be a trade-off between a need for generalization and retainment of the benefits of particulars. This is closely related to the problem of selecting landmarks.
- How do spatial reference descriptions depend on a task? Humans possess a repertoire of spatial reference systems. Furthermore, ad-hoc reference systems can be built. The question is whether it is possible to define a given task in a way such that we can decide

² <http://en.wikipedia.org/wiki/Orienteering>

which reference system should be selected from the repertoire or generated ad-hoc, in a manner that it is optimally suited to the goal. Also, accessibility to reference systems needs to be considered, that is: what reference systems are available to the audience? Goal and accessibility might be conflicting, e.g. the goal might be to give the most definite location, however the audience might not know these terms.

- Which action hierarchies need to be considered on different levels of detail and how do these determine spatial descriptions? An example would be the linkage between an event in time (e.g. “breakfast”) and a place where it happens (e.g. “downstairs”), and how a hierarchy is affected when a change occurs in terms of actions or goals.
- How can maps be reconstructed from narratives? For example: “There were 8 people sitting around the table, six with laptops in front of them and two with coffee.” The challenge is that it is possible to reconstruct multiple depictions from a single narrative such as the one above.
- Is there an ontology design pattern for spatial reference, along the ideas of S. C. Levinson? It would need to include a minimal set of spatial concepts (axis, relations) which are needed in order to describe relative, absolute and intrinsic reference systems. It would allow us to describe and publish robot experiments as well as human referencing on the Web.
- How can spatial reasoning be brought into the Semantic Web? Which reasoners, formalisms and spatial relations should be supported, and how can existing systems be integrated? Technically, SPARQL basic graph patterns and filters can be hacked in order to support existing spatial reasoners.
- How could we develop a computational model for answering “where” questions? An example of a “where” question is “Where is downstairs?” The question implies that communicating agents know what is meant with downstairs. The discussion led to the use-case of navigation instructions across different modes of transportation and different spaces, where the transportation modes might be driving, walking, cycling. Different modes of communication of the instructions should be considered (e.g. verbal, tactile and visual).

5.2 The outcomes

As a concrete outcome of the seminar, participants agreed to follow up on a number of promising interdisciplinary research topics. The latter include tasks such as paper writing, software development and formulations of technical challenges (with people in brackets being responsible for coordination):

1. Write a paper on a spatial reference design pattern. Automatic publication of descriptions of locations visited by a robot, using the Semantic Web as a sharing platform (Simon Scheider)
2. Offer qualitative spatial reasoning (e.g. RCC8++), as implemented in SparQ³, via a SPARQL engine (by result enrichment) based on the GeoSPARQL vocabulary⁴. Implement a software package (endpoint) which integrates spatial reasoners and could be

³ <http://www.sfbtr8.uni-bremen.de/project/r3/sparq/>

⁴ <http://www.opengeospatial.org/standards/geosparql>

- used with other reasoners, too. This idea already produced some results.⁵ (Thomas Scharrenbach)
3. Organize a hackathon on integrating spatial reasoning into the Semantic Web. (Brandon Bennett)
 4. Write a paper on grounding the Semantic Web. How do we propose to ground spatial terms in the Semantic Web? How do Luc's ideas, spatial science, cognitive science and the Semantic Web fit together? How to take into account the communication situation and the purpose of information? (Anne Cregan)
 5. Formulate a which and where question challenge: How to determine which reference frame is suitable for a certain where task, based on purpose and action? (Stephan Winter)
 6. Formulate a challenge for a Turing test for localization. Test whether robots and humans perform equivalently on localization tasks and locative description tasks (Diedrich Wolter)
 7. Formulate a challenge for building a gazetteer which interprets place descriptions. Examples for such descriptions are "the hill along Grattan street" or "in front of the door of the smithy in Port Khazard". Submitted as an InnoCentive challenge. (Werner Kuhn)

⁵ <https://qsr4sw.wordpress.com/>

Participants

- Brandon Bennett
University of Leeds, GB
- Mehul Bhatt
Universität Bremen, DE
- Saša Bodiroža
Humboldt-Univ. zu Berlin, DE
- Martin V. Butz
Universität Tübingen, DE
- Nicholas Chrisman
RMIT Univ. – Melbourne, AU
- Helen Couclelis
University of California – Santa Barbara, US
- Anne Cregan
Intersect – Sydney, AU
- Max J. Egenhofer
University of Maine, US
- Andrew U. Frank
TU Wien, AT
- Christian Freksa
Universität Bremen, DE
- Aldo Gangemi
CNR – Rome, IT
- Giancarlo Guizzardi
UFES – Vitoria, BR
- Verena V. Hafner
Humboldt-Univ. zu Berlin, DE
- Pascal Hitzler
Wright State University – Dayton, US
- Werner Kuhn
University of California – Santa Barbara, US
- Bruno Lara Guzman
Universidad Autónoma del Estado de Morelos, MX
- Vivien Mast
Universität Bremen, DE
- Mihoko Niitsuma
Chuo University, JP
- Kai-Florian Richter
Universität Zürich, CH
- Katharina J. Rohlfing
Universität Bielefeld, DE
- Yulia Sandamirskaya
Ruhr-Universität Bochum, DE
- Thomas Scharrenbach
Universität Zürich, CH
- Simon Scheider
Universität Münster, DE
- Michael Spranger
Sony – Tokyo, JP
- Stephan Winter
The University of Melbourne, AU
- Stefan Wölfl
Universität Freiburg, DE
- Diedrich Wolter
Universität Bamberg, DE

