Speaking of Location: Communicating about Space with Geospatial Natural Language

Kristin Stock
Massey Geoinformatics Collaboratory, Massey University, Auckland, New Zealand
k.stock@massey.ac.nz

Christopher B. Jones
School of Computer Science and Informatics, Cardiff University, Cardiff, United Kingdom
jonescb2@cardiff.ac.uk

Thora Tenbrink
School of Languages, Literatures and Linguistics, Bangor University, Bangor, United Kingdom
t.tenbrink@bangor.ac.uk

Abstract

Speaking of Location 2019 is the second edition of the Speaking of Location workshop series, which aims to foster transdisciplinary research to address the problem of automatic interpretation and generation of geospatial natural language. This introduction to the workshop proceedings provides background, discussing the definition and nature of geospatial natural language, presenting the papers contained in the proceedings volume, and situating them within the theoretical framework of The Semantic Pyramid, which is also described.

2012 ACM Subject Classification Artificial Intelligence → Natural Language Processing; Spatial-temporal systems → Geographic information systems

Keywords and phrases geospatial natural language, locative expressions, geographic information retrieval

1 Introduction

Speaking of Location 2019 was the second edition in the Speaking of Location workshop series, the first having been held in l'Aquila, Italy in 2017. Both workshops were held in conjunction with the long-running Conference on Spatial Information Theory. Speaking of Location workshops bring together researchers from across several disciplines, all of which have understanding that is needed to achieve the automatic generation, extraction and interpretation of natural language descriptions of geographic space. Linguistics researchers bring an understanding of the nature of language that describes location, and the way it is used in different cultures. Cognitive Science researchers bring an understanding of the ways that people conceptualise the world, particularly the geographic world, including the elements within it, and the way in which location is understood. Geography researchers bring an understanding of the geographic world itself, and the different ways in which it may be viewed, depending on purpose and environment. Computer Science researchers bring methods and theory that may be applied to automation of the complex and nuanced ways in which we describe location. All of these disciplines, and others, are required for the realisation of the vision of a system that can understand and generate human language about location in the way that humans do. Speaking of Location workshops are intended to facilitate, enable and encourage steps towards this vision.

In the following introduction to the workshop proceedings, we provide a brief outline of geospatial natural language and its complexities, before providing a conceptual framework (The Semantic Pyramid) and overview of the presented papers.

In Proceedings Speaking of Location 2019: Communicating about Space, Regensburg, Germany, September-2019.
Editors: K. Stock, C.B. Jones and T. Tenbrink (eds.);
Published at http://ceur-ws.org

2 What is Geospatial Language?

Humans routinely describe location using natural language. While locations are often described using place names, the use of terms and constructions that describe the location of an object relative to another is also common (e.g. *the church is opposite the bridge* or *the gardens are beside the Waikato River*). Such expressions can be very complex, containing multiple nested clauses and different reference frames. Thus location may be described relative to the observer, relative to another object (its location or its intrinsic properties) or relative to an external reference frame [9]. They may describe the location of an object as a static relationship to another object (as in the examples above), or dynamically using directions, focusing on how the location could be reached from some other location. In all cases, various perspectives can be used, such as the observer's perspective (static: *the shop is left of the church from my point of view*, dynamic: *you'll pass the shop on your left*), object property based (*the shop is in front of the castle entrance*), or compass-based (*the shop is north of the church*). Compass-based perspectives are often associated with survey perspectives, whereas observer's perspectives are commonly used in route descriptions [14].

A typical location description consists of three elements: the object whose location is being described (known as the locatum, figure or trajector), the object that is used as a reference (known as the relatum, ground, landmark or reference object), and the relationship between them, known as the spatial relation [9, 13, 8]. Spatial relations may take many forms [15], ranging from the very simple preposition (beside, opposite, near) to much more complex expressions in which verbs, adverbs, syntactic devices like commas or apostrophes and/or multi-word phrases also carry spatial information (for example, the road continues north beside the river, turning west where the river divides).

3 Communicating about Space across Different Contexts and Purposes

Context is a key element in the correct interpretation of natural language descriptions of location. Location descriptions may be made in different contexts, and for different purposes, and these have an impact on the way in which location descriptions are understood, and the degree of precision in our interpretation of the description. For example, the expression *the church is opposite the bridge* may be interpreted differently depending on the specific scene. Figure 1 shows two possible scenes that could be considered consistent with this expression, but that show quite different configurations. Thus if we don't have other details about the scene, there are often a number of different ways an expression could be interpreted, due in part to contextual factors, but also the tendency for natural language expressions to accommodate vagueness and underspecificity [1].

The way expressions are interpreted may also vary depending on what is possible or likely. For example, the preposition *on* has many senses [2], and we interpret the correct sense in part by considering what is the most likely, given the kinds of things we know about the way the environment is structured and what we are likely to encounter. For example, *Mary lives on the island* implies that Mary's house is physically on the surface of the island [5, 7, 3, 10], but this interpretation of the preposition on is not likely for the expression *Mary lives on the main street*, because houses are rarely built directly on the surface of the street itself. In contrast, we interpret this to mean that Mary's house is on the side of the main street, facing on to it. Conversely, this interpretation would not work for *Mary lives on the island*, as Mary would then live in the water (unless Mary lives in a house boat, but then we would

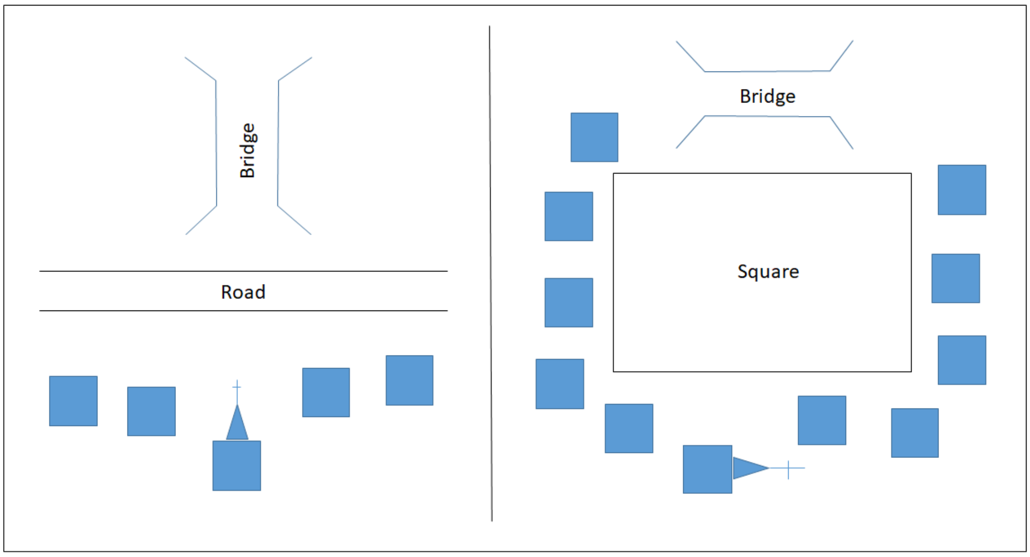

Figure 1 Example Interpretations of *the church is opposite the bridge*

be likely to use a different expression altogether).

4 The Semantic Pyramid

The well-known Semantic Triangle [11] (Figure 2) describes the relationship between Referents (real world objects and situations); Thoughts (human conceptualisations of them, which reflect the cognitive model that humans hold in their minds for those objects) and Symbols that we may use to represent them (e.g. natural/human language). This framework describes the relationship between objects and their cognitive and symbolic representations. As well as natural language, symbolic representations may include those in digital form, and while in the standard semantic triangle, these would be accommodated under the Symbols corner of the Triangle, here we present an expanded version of the Semantic Triangle that draws attention to the more specific details of digital representations of language and knowledge that are important for the task of interpretation and generation of geospatial natural language, and is illustrated in Figure 3. We create an additional triangular face, defined by two additional corners, one of which reflects the digital representation of natural language: the Language Representation corner. In order to work with geospatial natural language in computers, we transform it into a more structured form, extracting the elements that are important for the task of interpreting language. This may include the elements described in Section 2 (locata, relata, spatial relation), but may also include other related items like frame of reference. For example, [6]

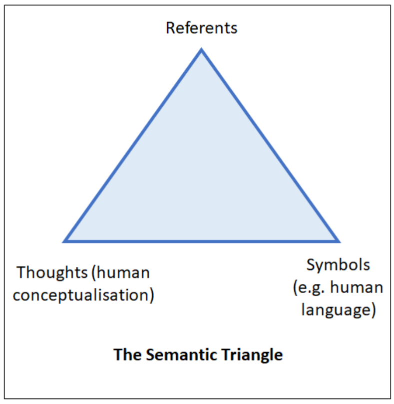

Figure 2 The Semantic Triangle

extract trajector, landmark and spatial indicator. The relationship between the Symbols (natural language) corner of the Semantic Pyramid (which formed a corner of the original Semantic Triangle) and the new, Language Representation corner, involves a process of abstraction, the former being the natural language that humans use, and the latter being a structured form of the pertinent elements that a machine can work with. The second new corner that appears in the Semantic Pyramid but was not in the Semantic Triangle, is that of knowledge representation, in which a computer stores a digital representation of a human conceptualisation. This might consist for example of a model of geographic space that encodes the way that entities in space relate to locations on earth and to each other as in the case of raster or vector representations, which reflect two different ways of conceptualising space; a domain ontology, which encodes the categories that communities use to describe the world and which reflect a common conceptualisation; and a logical formalism that supports reasoning about relations in geographic space (e.g. RCC8 [4]). Ontologies that are abstracted from natural language like GUM-Space [1] and ISO-Space [12] exist somewhere on the edge that joins the Language Representation corner to the Knowledge Representation corner, in that they vary in how closely they attempt to model human conceptualisations in contrast to their linguistic representation.

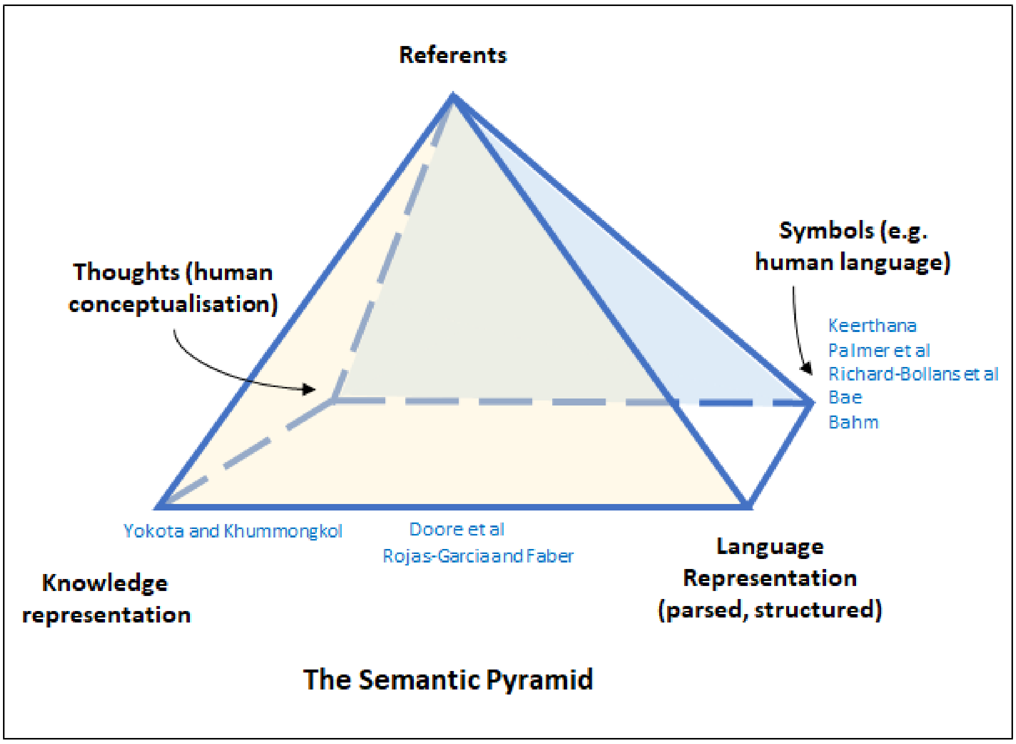

Figure 3 The Semantic Pyramid

In combination, the corners of the Semantic Pyramid reflect the range of research addressed by the Speaking of Location 2019 Workshop. We need to engage with these various fields of study in our efforts to adequately and comprehensively address the task of automated interpretation and generation of natural language location descriptions. Doing so requires input from a number of disciplines. In order to perform the interpretation and generation task, we need to understand how people conceptualise geographic objects, and this requires work in cognitive science and spatial cognition. We also need to understand the ways in which people use spatial language to describe the world, and particularly spatial scenes, and this requires the expertise of linguists. To identify, extract and represent relevant elements from language in a digital format, the expertise of linguists, computational linguists and computer scientists are required in combination. To identify, extract and represent relevant conceptual knowledge in a system to enable automated interpretation and generation, the expertise of computer scientists, working with cognitive scientists and geographers, is needed. Ultimately, the boundaries between these disciplines blur and hence a combined effort is required to address the challenge. This is the purpose of the Speaking of Location Workshop series. In 2019, in addition to a tutorial on Cognitive Discourse Analysis [16], position pitches, a keynote and a panel discussion, we included 8 research papers in the workshop and the proceedings. These papers occupied different locations on the corners and edges of the Semantic Pyramid.

Clustered around the Symbols corner of the pyramid are papers that focus on the nature and characteristics of spatial language, including **Keerthana's** paper, which describes the nature of boundedness in path descriptions in the Mayalayam language, and then looks at how this is related to paths that are stratified (consist of multiple segments); **Palmer, Blythe, Gaby, Hoffmann and Ponsonnet** who discusses the use of frame of reference in Australian Aboriginal languages, calling into question the common view that absolute, cardinal direction-focused frames of reference are dominant in these languages, and revealing a strong link to landscape. These papers focus on the nature and understanding of spatial language per se, rather than automation, but are useful for the project of interpreting and generating geospatial natural language in that they help us to understand the ways in which people use spatial language, and which we therefore need to consider when attempting to automate interpretation or generation of location descriptions. These cross-linguistic studies also help us to understand the range of variations that are possible between languages, understanding that is essential for the creation of automated systems that can be used in a multi-lingual environment.

Also focused on the understanding of spatial language (Symbols corner) is **Richard-Bollans, Gómez Álvarez, Bennett and Cohn**, which describes a tool that is designed to collect data to enable examination of the ways in which people understand spatial relations; and then provides analysis of data collected using that tool and **Bae**, who presents a paper that discusses the narratives that are used during joint route planning, applying a Conversation Analytic framework to study the structure of the interactions. **Bahm** combines a focus on wayfinding with that on a specific spatial relation: *through*, identifying three types of relationships between the *through* preposition and the scene within which it is applied, in a museum environment. Again, the focus of this work is at the Symbols corner of the pyramid. Moving around the base of the pyramid to the Language Representation corner, **Rojas-Garcia and Faber** describe a semi-automatically extracted representation of geographic terms from a text corpus, specifically describing rivers and bays, from which they create semantic networks to explore underlying knowledge, thus moving towards the Knowledge Representation corner. Addressing both the Language Representation and Knowledge Representation corners, **Doore, Sarrazin and Giudice** present a place graph model containing key, abstracted aspects of natural language descriptions in order to support a user interface for navigation by blind and visually impaired museum visitors. Finally, at the Knowledge Representation corner of the semantic pyramid, and the edges that connect it to the Concepts and the Language Representation corners, **Yokota and Khummongkol** apply

Mental Image Directed Semantic Theory to the problem of natural language understanding in the robotics domain, linking machine and human language and conceptualisation. As can be seen in Figure 3, the largest number of papers are at the Symbols corner of the pyramid, addressing the nature of spatial and/or geospatial language, with few papers at the other corners. While the study of spatial language is long standing, the incorporation of this work into the multidisciplinary environment typified by the Speaking of Location workshop series may reflect recognition of the need to consider the complexities of spatial language as this field of research (addressing the automation of location description interpretation and generation) matures. In particular we can regard this transdisciplinary approach as part of a process of moving beyond what might be regarded as a superficial understanding to the formulation and implementation of computable models of the way in which we describe location using natural language.

5 Conclusions

The papers contained in these proceedings, and their presentation and discussion at the Speaking of Location 2019 workshop provide a step towards the goal of realising the vision of automated interpretation and generation of geospatial natural language, with all its complexities, vagueness, under-specificity, context sensitivity and dynamism. The challenges are still significant, and the research opportunities numerous, but the benefits of achieving such a vision in the current text-rich environment indicate that this vision will continue to be important going into the future.

References

1 John A. Bateman, Joana Hois, Robert Ross, and Thora Tenbrink. A Linguistic Ontology of Space for Natural Language Processing. *Artif. Intell.*, 174(14):1027–1071, September 2010. URL: `http://dx.doi.org/10.1016/j.artint.2010.05.008`, `doi:10.1016/j.artint.2010.05.008`.

2 Kenny R. Coventry, Richard Carmichael, and Simon C. Garrod. Spatial Prepositions, Object-Specific Function, and Task Requirements. *Journal of Semantics*, 11(4):289–309, January 1994. URL: `https://academic.oup.com/jos/article/11/4/289/1640110`, `doi:10.1093/jos/11.4.289`.

3 Andrew Frank and Martin Raubal. Formal specification of image schemata – a step towards interoperability in geographic information systems. *Spatial Cognition & Computation*, 1(1):67–101, 1999. URL: `insights.ovid.com`.

4 Nicholas M Gotts, John M Gooday, and Anthony G Cohn. A connection based approach to common-sense topological description and reasoning. *The Monist*, 79(1):51–75, 1996.

5 Mark Johnson. *The body in the mind: The bodily basis of meaning, imagination, and reason.* The body in the mind: The bodily basis of meaning, imagination, and reason. University of Chicago Press, Chicago, IL, US, 1987.

6 Parisa Kordjamshidi, Martijn Van Otterlo, and Marie-Francine Moens. Spatial Role Labeling: Towards Extraction of Spatial Relations from Natural Language. *ACM Trans. Speech Lang. Process.*, 8(3):4:1–4:36, December 2011. URL: `http://doi.acm.org/10.1145/2050104.2050105`, `doi:10.1145/2050104.2050105`.

7 George Lakoff. *Women, Fire, and Dangerous Things - What Categories Reveal about the Mind.* The University of Chicago Press, 1987.

8 Ronald W. Langacker. An Introduction to Cognitive Grammar. *Cognitive Science*, 10(1):1–40, 1986. URL: `https://onlinelibrary.wiley.com/doi/abs/10.1207/s15516709cog1001_1`, `doi:10.1207/s15516709cog1001_1`.

9 Stephen C Levinson. *Space in language and cognition: Explorations in cognitive diversity*, volume 5. Cambridge University Press, 2003.

10 David M Mark and Andrew U Frank. Experiential and formal models of geographic space. *Environment and Planning B: Planning and Design*, 23(1):3–24, 1996.

11 C.K. Ogden and I.A. Richards. *The meaning of meaning*. Harcourt Brace Jovanovich, New York & London, 1923.

12 James Pustejovsky. ISO-Space: Annotating Static and Dynamic Spatial Information. In Nancy Ide and James Pustejovsky, editors, *Handbook of Linguistic Annotation*, pages 989–1024. Springer Netherlands, Dordrecht, 2017. URL: `https://doi.org/10.1007/978-94-024-0881-2_37`, `doi:10.1007/978-94-024-0881-2_37`.

13 Leonard Talmy. How Language Structures Space. In Herbert L. Pick and Linda P. Acredolo, editors, *Spatial Orientation: Theory, Research, and Application*, pages 225–282. Springer US, Boston, MA, 1983. URL: `https://doi.org/10.1007/978-1-4615-9325-6_11`, `doi:10.1007/978-1-4615-9325-6_11`.

14 Holly A. Taylor and Barbara Tversky. Perspective in Spatial Descriptions. *Journal of Memory and Language*, 35(3):371–391, June 1996. URL: `http://www.sciencedirect.com/science/article/pii/S0749596X96900212`, `doi:10.1006/jmla.1996.0021`.

15 Thora Tenbrink. *Space, Time, and the Use of Language : An Investigation of Relationships*. Mouton de Gruyter, Berlin, 2007.

16 Thora Tenbrink. *Cognitive Discourse Analysis: An Introduction*. Cambridge University Press, in press.