



Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Winter, S;Baldwin, T;Tomko, M;Renz, J;Kuhn, W;Vasardani, M

Title:

Spatial concepts in the conversation with a computer

Date:

2021-06-21

Citation:

Winter, S., Baldwin, T., Tomko, M., Renz, J., Kuhn, W. & Vasardani, M. (2021). Spatial concepts in the conversation with a computer. Communications of the ACM, 64 (7), pp.82-88. <https://doi.org/10.1145/3440752>.

Persistent Link:

<https://hdl.handle.net/11343/307561>

Spatial concepts in the conversation with a computer

STEPHAN WINTER, TIMOTHY BALDWIN, and MARTIN TOMKO, The University of Melbourne, Australia

JOCHEN RENZ, Australian National University, Australia

WERNER KUHN, University of California at Santa Barbara, USA

MARIA VASARDANI, RMIT University, Australia

Conversing about places with a computer poses a range of challenges to current AI. The content of the communication strongly depends on the personal perspectives and preferences of the individuals engaged in the conversation, and may contain high levels of ambiguity that can only be resolved via mutual information and local context. Among the challenges to converse about places are the broader conceptual nature of places across geographic scales, the flexibility, vagueness and ambiguity of language, the context-dependency impacting on language grounding and disambiguation, and the interpretation of *part-of* relations and other spatial relations. This review illustrates the challenges, and summarizes emerging work to enable computers to become a conversational partner of humans on places and locations.

CCS Concepts: • **Information systems** → **Location based services**; • **Human-centered computing** → **Natural language interfaces**.

Additional Key Words and Phrases: geographic information, space-place dichotomy, voice interface

ACM Reference Format:

Stephan Winter, Timothy Baldwin, Martin Tomko, Jochen Renz, Werner Kuhn, and Maria Vasardani. 2018. Spatial concepts in the conversation with a computer. 1, 1 (June 2018), 11 pages. <https://doi.org/10.1145/1122445.1122456>

KEY INSIGHTS

- Conversing about places challenges machine-understanding due to the flexibility, ambiguity and context-dependency of place references, and the vagueness of the nature of places.
- Conversing about places assumes an understanding of the structural and hierarchical knowledge of the world in order to use and interpret qualitative spatial prepositions.
- Yet, discourse patterns about places can be modelled or learned, and translated into formal interactive conversation patterns, thus supporting the development of intelligent assistance systems.

Authors' addresses: Stephan Winter, winter@unimelb.edu.au; Timothy Baldwin, tbaldwin@unimelb.edu.au; Martin Tomko, tomkom@unimelb.edu.au, The University of Melbourne, Parkville, VIC, 3010, Australia; Jochen Renz, jochen.renz@anu.edu.au, Australian National University, Canberra, ACT, 2600, Australia; Werner Kuhn, werner@ucsb.edu, University of California at Santa Barbara, Santa Barbara, CA, USA; Maria Vasardani, mvasardani@rmit.edu.au, RMIT University, Melbourne, VIC, 3000, Australia.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2018 Association for Computing Machinery.

Manuscript submitted to ACM

Manuscript submitted to ACM

Illustrative scenario

Kim, who is shopping for comfortable shoes, turns to her intelligent conversational agent Marvin

Kim, looking at an advertisement: *Where can I buy these sneakers? It needs to be a physical store because I'm not sure of my size in that brand.*

Marvin *Hype in the Imperial Shopping Centre, in the West of the city, currently have the best price.*

Kim *OK, let's drive there.* She connects Marvin to a self-driving vehicle.

...

Marvin *The main car park is full.*

Kim *OK, then take me to the rear one, behind the cinemas. But I've never entered the shopping center from the back.*

Marvin *There are elevators at the cinema. Hype is on the second floor. It is near the atrium, where Pens&Paper used to be.*

...

Marvin *The sneakers should be at the back of the store to the right, in the 'casual wear' section.*

Compare this conversation to the maps or to the step-by-step route descriptions provided by current location-based services: this conversation is concise, situated, and personalized. More importantly, it is a conversation about 'where' that does not rely on well-known toponyms, points-of-interest, or on coordinates. It uses *places* and *their relations*. This is precisely what current AI is not capable of.

1 INTRODUCTION

Human interactions with the physical environment are often mediated through information services, and sometimes depend on them. These human interactions with their environment relate to a range of scales [28], in the scenario above from the "West of the city" to the "back of the store", or beyond the scenario to "the cat is under the sofa". These interactions go far beyond references to places that are recorded in geographic gazetteers [37], both in scale (the place where the cat is) and conceptualization (the place that forms the West of the city [29]), or that fit to the classical coordinate-based representations of digital maps. And yet, these kinds of services have to use such digital representations of environments, such as digital maps, building information models, knowledge bases, or just text/documents. Also, their abilities to interact are limited to either fusing with the environment [44], or using media such as maps, photos, augmented reality, or voice. These interactions also happen in a vast range of real-world contexts, or *in situ*, in which conversation partners typically adapt their conversational strategies to their interlocutor, based on mutual information, activities and the shared situation [2]. Verbal information sharing and conversations about places may also be more suitable when visual communication through maps or imagery is inaccessible, distracting, or irrelevant, such as when navigating in a familiar shopping mall.

Much of human conversation is about places without well-known names or with only local significance, such as *the cinemas* (see Scenario). They are often identified not by names at all, but by qualities (*the rear car park*; more examples in Table 1). Sometimes, places lack a well-established location or extent, such as *the West of the city*, or are characterized in their location just by spatial relations (*the elevators at the cinema*).

Thus, we highlight here to broader concept of *place* we are confronted with. Borrowing from the *core concepts* [24] we say that every object (in the core concepts) is located. And places are objects [32]. To be more precise, references to objects make them places if the conversation is about localization – in contrast to other object properties. But

Example	Places referred to through ...	Scale
<i>under the sofa</i>	a qualitative spatial preposition with a located object	hyperlocal
<i>This place to sit down</i>	an affordance facet	hyperlocal
<i>The beautiful beach</i>	a quality facet and a generic noun	local
<i>The West of the city</i>	a vague and ambiguous place description	city
<i>The capital of Slovakia</i>	through a function of a place within another gazetted place	city
<i>The Newer Volcanic Province</i>	a vague region with a gazetted placename	regional
<i>The river crossing Austria and Slovakia</i>	a qualitative spatial relation with gazetted placenames	country

Table 1. Illustrations of additional complexities in references to places.

localization of an object can happen in a variety of ways, and geometric localization – i.e., in a spatial reference frame – is only one of them. Another one is the one by qualitative spatial relations. In this article we focus on the latter.

But while vital for human information sharing, decision-making, and coordination, conversations about places are still poorly supported by information services and situated robotic assistants such as self-driving vehicles, or, in our scenario, by Marvin. Typical conversational contexts about places include:

- Partners establishing a shared situation awareness, such as in command and control [10] for emergency response (*Point the firehose more to the right*), in defence (*The sniper must be on the roof of the building next to the tree*), and in collaborative scenarios in games.
- Conversational location-aware services used as recommender systems (*Where can I buy these sneakers?*) or self-driving vehicles (*Go to the car park behind the cinemas*).
- Services that obfuscate or abandon geometric descriptions in favor of vague and qualitative descriptions to protect private or sensitive information (*Kim lives in the city's East* is protecting the exact home address by referring to a vague and coarser place).

The wide-ranging opportunities to engage users in a verbal information exchange when seeking and providing place information extend beyond pure information seeking, and include commands as well the provision of explanations to machines. Here, we outline the underlying challenges of conversing about places with computers, and report on recent progress in this space. We also identify salient areas for immediate and longer-term future work.

2 SCOPE AND CHALLENGES

Communication about our physical environment – small scale and large scale – is special and difficult: the content of the communication strongly depends on the personal perspectives and preferences of the individuals engaged in the conversation, and may contain high levels of ambiguity that can only be resolved via mutual information and local context. Yet, this communication needs to enable physical interaction with the environment and others in it, such as coordination or planning of movement [5]. In conversations with a computer, references to both places and their relations are difficult to parse and represent formally for a variety of reasons, among them the inherent flexibility, vagueness and ambiguity of language (Table 1), the context-dependency impacting on language grounding and disambiguation, and the interpretation of *part-of* relations (scale) and other spatial relations.

2.1 Language and facets of place

When conversing about places, the expressions used in linguistic strategies convey place qualities (*the beautiful beach*), relations between places (*the river crossing Austria and Slovakia*) and relations between people and places, or roles of places. Such functional, social, and operational roles of places are imposed by individuals, or groups. A function of a place can be tied to what it affords to an individual. The social role is derived from ownership. And the operational role is derived from intention (and spatial configuration). A recent comprehensive review [15], using facet theory [3], categorised these means of expression into three top-level categories of facets of places. *Anthropocentric* facets relate to the relations between people and places. They include the affordances of places and activities linked to them, as well as emotive facets that cover the sense of place and cultural attachment. *Geographic* facets describe the spatial and physical properties of places. Spatial facets enable localization of a place in space and describe the meaning of the place in relation to space and other places, while physical facets capture the structure, material, or form of a place. And finally, *linguistic* facets present a distinct category, capturing purely linguistic manifestations of place through place-names and identifiers. Linguistic facets can pick up anthropocentric facets (*the sneaker store*) or geographic facets (*the Lower Eastside*), but do not have to (*Hype*).

While the bulk of research and development in assistance systems thus far has focused only on linguistic facets (e.g., *place names*), it falls short in the interpretation of references to place in natural language conversations. Situated assistants that are able to interact with the physical environment such as self-driving vehicles will require richer support for multi-faceted natural language interaction about places. In this direction points already work in controlling robots by spatial references, for example identifying objects in an environment that can be referred to on the basis of a range of spatial reference systems [30, 41], or reflecting on deductive human reasoning with preferred models instead of formal qualitative spatial reasoning [34].

2.2 Vagueness

References to places are inherently vague in their specification of a location [32]. Since place descriptions are usually *referring expressions* – expressions *uniquely* identifying a referent in a set of possibilities – they are typically as specific as necessary [40]. This specificity is given not by coordinates – i.e., neither by crisp nor by fuzzy geometric descriptors (e.g., [8]) –, but by contrast [45]. For example, the location of *the West of the city* [29, 33] is sufficiently described as in contrast to the East, South, or North, and *the back of the store* is sufficiently specified in contrast to its front. Vagueness also applies to the qualities of places (*the full car park*), and to the relations between places (*elevators at the cinema*), each with an intent to be only as distinctive as necessary to be relevant in the given context.

2.3 Ambiguity

Most references to places are ambiguous. While, for example, country names are globally unique, the same is already no longer true for city names (*Paris*) or landscape names (*Alps*). *The sneakers shop on the second floor* is not a globally unique reference. Disambiguation requires further locative discriminators, which are enabling the popular cluster-based disambiguation, further non-locative discriminators, which are enabling iterative identification such as through dialogue, and references to spatial relations, which are enabling disambiguation by spatial reasoning. Popular spatial relations are containment hierarchies, emphasizing the importance of scale (*the sneakers shop in the Imperial Shopping Centre [in this city, in this country]*), but may also be a relation with a local landmark (*the sneakers shop near the food court*). Dialog-based disambiguation methods [46] still require intense research, as they flip the roles of the computer and the

human user, in that the human user becomes the source of information provided to the computer. Verbal disambiguation by all three approaches is easier for a situated than for a non-situated conversational agent.

2.4 Scale

Scale, or containment, is an intrinsic property of places that is impacting on spatial reasoning and enabling to adapt to spatial context. The question *Where can I buy those sneakers* can be answered correctly but of varying relevance by *in the West of the city* (e.g., telling Kim the general direction), *at the shopping mall* (e.g., if Kim is in the neighborhood), *on the second floor* (if Kim is in the car park), or by *in the back of the shop* (if Kim is entering the shop). These descriptions range between scales, and benefit from containment hierarchies that allow to zoom and reason between levels of granularity [36].

2.5 Context

Conversational context provides implicit information taken into account when generating or interpreting place references. While this elementary function of context is well understood and undisputed, the concept of context itself is still ill-defined [9]. Context is broader than the immediate perceptual cues of a situated agent, such as positioning [39], orientation and vista analysis for meaningful salient features. It includes a consideration of the human communication partner, of the physical environment beyond the vista, and of social and technical aspects [19]. Despite of this common broader understanding, and in tacit recognition of the ill-defined challenge, existing research on context-aware computing often links context only to position or any semantically richer notion of location. For example, some work matches a position with points-of-interest and the activities they afford [e.g., 16, 38], and on the smaller scales of smart environments sensors are used for activity recognition, which links to location [1], or a broader notion of location is introduced to deal with positioning uncertainty in these smart environments [4]. Semantically richer notions of location can enable assuming semantically meaningful places, including their relative description (2.1) by selecting *salient* facets [36]. It also includes the choice of a degree of specificity that is *sufficient* for the communication purpose (2.2). It requires selection of *relevant* facets to avoid unnecessary ambiguity (2.3). And it requires a choice of scale that *relates* to the current decision making (2.4). Human conversation partners naturally adapt to context even when they are physically removed (e.g., speaking on the phone, a form of co-presence [47]).

Descriptions can be personalized for familiarity (in contrast to Kim, a newcomer to the shopping mall would also need instructions on how to get to the second floor); ability (a person in a wheelchair might require nuanced instructions on how to access the second floor); role (a cleaner might have access to staff-only entrances); and adapted to environmental characteristics (a complex environment will generally require a more detailed description) and dynamics (think of the full car park). This context-dependency in question-answering is currently avoided in most route direction services that maximally enrich instructions by a few personalized landmarks selected from the user's history.

Similarly, the nearness relation is known to be context-dependent and occasionally even asymmetric. For example, the *sneaker shop near the food court* and the *airport near the city centre* require different handling of the proximal relation. The interpretation of proximity prepositions is a persisting challenge.

3 ANATOMY OF A PLACE ASSISTANCE SYSTEM

In order to converse about place, an assistance system needs to manage the challenges in handling place information along the four interleaved tasks of capturing an individual's expression as it refers to places (covering the parsing of

natural language), modelling extracted knowledge about places in some representation, processing such knowledge, and, thus, finally, interacting with the individual.

3.1 Capture

Place knowledge expressed in common language – questions or answers, commands, or explanations – can be parsed, and place references extracted and interpreted within their conversational context. This necessarily includes the ability to parse and interpret the rich means of referring to the various facets of place, including but not limited to toponym resolution and disambiguation.

The intricacies of handling context, scale, and vagueness become salient when locative expressions from natural language interactions are matched against the content of a knowledge base, with the main challenge of resolving ambiguities. When these knowledge bases are geometric ones (*spatial databases*) this matching is called geocoding [12, 13, 27]. The knowledge base can also be a qualitative one; see next section.

Some challenges in this matching process concern the unification between different references to the same place (e.g., finding out that *the sneakers shop* and *Hype* are actually one and the same place in particular contexts), and the disambiguation where the same term refers to different places (e.g., finding out that *Hype* in one conversation is a different place to a *Hype* store in another context). A related challenge is to establish that such a reference refers to an object missing in the knowledge base (e.g., if the mall is in the knowledge base, but no *Hype*), or to a changed one (e.g., if *Hype* is referring to the place that is stored in the knowledge base still as *Pens&Paper*) [6, 33, 35].

3.2 Modeling

Place knowledge is inherently relational, structured by encoding primarily qualitative relations between places encoded in place references, instead of coordinate-based grounding. The representation of such relational place knowledge – whether inherent in the questions, captured in knowledge bases or encoded in answers – requires relational data structures adapted for place reasoning. Such a representation is by so-called *place graphs* [43] – see Figure 1 for an example –, an application of the concept of *property graphs*. Property graphs have recently received intensive attention in the database community, and are the basis of the first new query language standard considered by ISO in 35 years (<https://www.gqlstandards.org/resources/committees-and-processes>). Place graphs provide also the means to support the capture of conversational context, scale and rich facets of place [7].

Place graph structures enable the capture of knowledge about relations between places even in cases where places have vague extents. As place graphs do not need to be anchored into a geographical coordinate reference system, the knowledge captured enables the reasoning flexibility observed when humans reason about places [6]. In particular, the matching of the place knowledge against the knowledge gathered from additional utterances during conversations enables local consistency checking, without requirements for perfect matching against a globally consistent place knowledge base.

3.3 Reasoning

Reasoning refers to the task of inferring unknown information from known or given information. Since natural language uses predominantly qualitative relations, the preferred reasoning is qualitative reasoning. For example, if *the sneakers shop* is *inside the shopping center*, and the *shoes* are *inside the sneakers shop*, it is possible to infer with certainty that the shoes are also inside the shopping center. This kind of compositional transitive reasoning relies on formalisms that precisely define vague spatial relations between entities [26].




Fig. 1. A graph of the place knowledge extracted from the conversation in the scenario.

Yet, individuals may have their own vague conceptualizations of common spatial relations [22, 23], and a richer set of spatial prepositions than a formal model of relations provides [30]. They also can flexibly choose spatial reference frames [21, 25, 41]. For example, telling Kim *to the right* can refer to the walking direction, the actual body orientation, or the shop layout. The interpretation of those projective prepositions requires access to viewpoints, and hence a situated assistant. Its reasoning has to carry the reference frame explicitly, and the determination whether new information is consistent with previous information remains a challenging problem.

While reasoning with predefined sets of relations with well-described compositions is relatively simple, reasoning between *different* types of spatial relations (e.g., a direction relation and a distance relation) is still an open problem. Furthermore, in certain place-related assistance tasks, the standard way of doing spatial reasoning is not applicable. For example, if the food court is left from the elevator, and Hype is left from the food court, for formal reasoning it remains ambiguous what the relation between Hype and the elevator is (*left* or *back*). But even if it were known, the composition is typically not helpful since Kim has to pass the food court first in this example and cannot walk directly in the direction of Hype. Instead, a valid inference needs to concatenate the two paths. In order to deal with the problem of context and user location, we have developed a unified representation of direction relations that allows us to disambiguate relative direction terms and to reason about them uniformly [18].

3.4 User Interaction

The initial – and still dominant – focus of place assistance systems about places [20]. This is now mostly happening ubiquitously and pervasively, including through conversational assistants such as Cortana and Siri. Yet, with the nearing commodification of autonomous systems, users will no longer only serve as consumers of knowledge served by computers, but will increasingly also communicate knowledge to computers that will act on it, including statements about preferences and refined specifications (*Get me to the shopping center – Which entrance would you prefer?*).

This transition to bi-directional interaction will pose new challenges to context handling and language grounding, personalization and privacy [42], and consideration of the physical affordances of the places as experienced by the combined user-machine system (*The rear entrance, please – Sorry, we lack after-hours access to the rear car park*). Also, language grounding and question-answering systems have thus far focused only on supporting the verbal interaction of users and machines either at the hyperlocal scale (vista space) [5, 30], or conversely, the global geographical scale with relatively coarse place knowledge [31], while place assistance as in our scenario interacts across all scales in an integrated manner. Furthermore, while situated interaction offers significant contextual clues for refined interpretations of questions, the modalities for machines to capture the situation and to seek additional feedback from users are still in their infancy. This is related to the ability to converse in dialogue as well as to capture facets of the environment that go beyond location and personalization of the user. In particular, the ability of machines to reflect on vague and multifaceted place knowledge is thus far weak both in construction (*West of the city*) and in reasoning (*behind the cinemas*).

4 A ROADMAP

The challenges of conversing about place have not been resolved by artificial intelligence so far, neither conceptually nor technologically. This is primarily caused by:

- Spatial databases (aka geographic information systems) remain strongly grounded in geometry and therefore lack capabilities to handle either the above challenges of capturing information about places (instead enforcing a crisp, scale-dependent, geometry-grounded spatial extent), or for computing and reasoning with qualitative spatial relations.
- Qualitative spatial representations and reasoning [26] capture some of the qualitative spatial relations referred to in language, but in formal interpretations that are not designed to directly cope with the flexibility of language and the embedding of the relations in conversational context. For example, we can reason with *near* yet we lack the ability to ground its meaning.
- Machine learning methods are neither suited to disambiguate places nor to interpret spatial relations. Due to the cross-cutting factors described above (context-dependence, scale, facets, vagueness, and ambiguity), statistical models often fail to detect significant correlations, or conversely, find spurious correlations – the challenge to interpret the meaning of *near* remains an excellent example.

In this situation, the anatomy of an idealised place-based assistance system as sketched in our scenario is intended to inspire a vision of a true conversational interaction with a computational assistant that can act in the physical environment, and interact with the user in a natural manner. This vision also enabled us to review some of the progress towards this vision, and to demonstrate that the state-of-the-art approaches are neither sufficient nor have we sufficiently tried to integrate the wide field of research involved.

A first, and enabling stepping stone is the establishment of a knowledge base of place knowledge. The obvious choice for a gold-standard of place knowledge is human verbal expressions – they are qualitative, contextualized (hence, sufficiently disambiguated), and independent from geometric descriptions, from commonly agreed collections of places such as gazetteers, or from named entities in general. Instead, the knowledge base is set up by places committing to an ontology of objects used to refer to (abstract) locations. Note that others have pointed to the contested nature of such extraction [11].

The next step is extending qualitative spatial representations and reasoning models to enable the handling of context. For example, Hua et al. [17] have shown that reasoning with relative directions (*left*, *right*) involves a non-trivial ternary relation of two explicit referents and a third, possibly implicit anchor point which has to be identified from the context of the conversation.

A third critical step towards the vision presented here is the ability to effectively query place knowledge bases in natural language. This step profits from learning human answering patterns on place-related questions [14]. The translation of natural language questions into queries (e.g., GeoSPARQL) can already be achieved. Thus the only unresolved dependency for this last step is knowledge bases that contain place knowledge including qualitative spatial relations, which cannot be computed from geometry without context.

More importantly even, this paper makes a case that it is the integration of these areas that will be required to make real progress. In that sense we encourage work on the intersection between spatial databases and machine learning, or natural language processing, in the construction of knowledge bases such as place graphs, and then further between spatial databases, machine learning, and qualitative spatial reasoning for approaching contextualized reasoning.

In summary, situated conversations about places require the machine to understand verbal place-based expressions (commands, descriptions, or questions) and to respond adequately, accessing place knowledge bases. We will soon no longer just ask conversational assistants about well-known place knowledge (*What is the capital of Slovakia?*), but will require systems to understand complex interactions about places, enable machines to seek additional specifications to resolve uncertainty or ambiguity, and have the facility to add new facts extracted from such conversations to its knowledge base. To achieve this, however, a concerted research effort in spatial information science, natural language processing, database management, qualitative spatial reasoning, and AI is necessary.

5 ACKNOWLEDGMENTS

This work has been supported by a Discovery Project funded by the Australian Research Council (Grant DP170100109).

REFERENCES

- [1] Marjan Alirezaie, Jennifer Renoux, Uwe KÄckemann, Annica Kristoffersson, LarsÄ Karlsson, Eva Blomqvist, Nicolas Tsiftes, Thiemo Voigt, and Amy Loutfi. 2017. An Ontology-based Context-aware System for Smart Homes: E-care@home. *Sensors* 17, 7 (2017), 1586:1–23. <https://www.mdpi.com/1424-8220/17/7/1586>
- [2] Andrea Ballatore and Michela Bertolotto. 2015. Personalizing maps. *Commun. ACM* 58, 12 (2015), 68–74. <https://doi.org/10.1145/2756546>
- [3] David Canter. 1997. The Facets of Place. In *Toward the Integration of Theory, Methods, Research, and Utilization*, Gary T. Moore and Robert W. Marans (Eds.). Springer US, Boston, MA, 109–147. https://doi.org/10.1007/978-1-4757-4425-5_4
- [4] Pedro Chahuara, FranÄois Portet, and Michel Vacher. 2017. Context-aware decision making under uncertainty for voice-based control of smart home. *Expert Systems with Applications* 75 (2017), 63–79. <https://doi.org/10.1016/j.eswa.2017.01.014>
- [5] Joyce Y. Chai, Rui Fang, Changsong Liu, and Lanbo She. 2017. Collaborative Language Grounding Toward Situated Human-Robot Dialogue. *AI Magazine* 37, 4 (2017), 32–45. <https://doi.org/10.1609/aimag.v37i4.2684>
- [6] Hao Chen, Maria Vasardani, and Stephan Winter. 2018. Georeferencing places from collective human descriptions using place graphs. *Journal of Spatial Information Science* 2018, 17 (2018), 31–62. <https://doi.org/10.5311/JOSIS.2018.17.417>
- [7] Hao Chen, Maria Vasardani, Stephan Winter, and Martin Tomko. 2018. A Graph Database Model for Knowledge Extracted from Place Descriptions. *International Journal of Geo-Information* 7, 6 (2018), Paper 221 (30 pages). <https://doi.org/10.3390/ijgi7060221>
- [8] C. Derungs and Ross Purves. 2014. From text to landscape: Locating, identifying and mapping the use of landscape features in a Swiss Alpine corpus. *International Journal of Geographical Information Science* 28, 6 (2014), 1272–1293.
- [9] Anind K. Dey, Gregory D. Abowd, and Daniel Salber. 2001. A Conceptual Framework and a Toolkit for Supporting the Rapid Prototyping of Context-Aware Applications. *Human-Computer Interaction* 16, 2 (Dec. 2001), 97–166. https://doi.org/10.1207/S15327051HCI16234_02
- [10] Mica R. Endsley and Debra G. Jones. 2011. *Designing for Situation Awareness: An Approach to Human-Centered Design*. CRC Press, Boca Raton, FL.
- [11] Heather Ford and Mark Graham. 2016. Provenance, power and place: Linked data and opaque digital geographies. *Environment and Planning D: Society and Space* 34, 6 (2016), 957–970.

- [12] Daniel W. Goldberg, John P. Wilson, and Craig A. Knoblock. 2007. From Text to Geographic Coordinates: The Current State of Geocoding. *URISA Journal* 19, 1 (2007), 33–46.
- [13] Milan Gritta, Mohammad Taher Pilehvar, and Nigel Collier. 2018. Which Melbourne? Augmenting Geocoding with Maps. In *56th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Melbourne, Australia, 1285–1296. <https://doi.org/10.18653/v1/P18-1119>
- [14] Ehsan Hamzei, Haonan Li, Maria Vasardani, Timothy Baldwin, Stephan Winter, and Martin Tomko. 2019. Place questions and human-generated answers: A data analysis approach. In *Geospatial Technologies for Local and Regional Development*, Phaedon Kyriakidis, Diofantos Hadjimitsis, Dimitrios Skarlatos, and Ali Mansourian (Eds.). Springer, Cham, Switzerland, 3–19.
- [15] Ehsan Hamzei, Stephan Winter, and Martin Tomko. 2020. Place facets: a systematic literature review. *Spatial Cognition & Computation* 20, 1 (2020), 33–81. <https://doi.org/10.1080/13875868.2019.1688332>
- [16] Eric Horvitz, Carl Kadie, Tim Paek, and David Hovel. 2003. Models of Attention in Computing and Communication: From Principles to Applications. *Commun. ACM* 46, 3 (March 2003), 52–59. <https://doi.org/10.1145/636772.636798>
- [17] Hua Hua, Jochen Renz, and Xiaoyu Ge. 2018. Qualitative Representation and Reasoning over Direction Relations across Different Frames of Reference. In *16th International Conference on Principles of Knowledge Representation and Reasoning*. AAAI, Arizona State University, USA, 551–560.
- [18] Hua Hua, Peng Zhang, and Jochen Renz. 2019. Qualitative Place Maps for Landmark-Based Localization and Navigation in GPS-Denied Environments. In *27th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. Association for Computing Machinery, Chicago, IL, USA, 23–32. <https://doi.org/10.1145/3347146.3359107>
- [19] Haosheng Huang, Georg Gartner, Jukka M. Krisp, Martin Raubal, and Nico Van de Weghe. 2018. Location based services: ongoing evolution and research agenda. *Journal of Location Based Services* 12, 2 (2018), 63–93. <https://doi.org/10.1080/17489725.2018.1508763>
- [20] Christopher B Jones and Ross S Purves. 2009. Geographical information retrieval. In *Encyclopedia of Database Systems*. Springer, Boston, MA, 1227–1231.
- [21] Roberta L. Klatzky. 1998. Allocentric and Egocentric Spatial Representations: Definitions, Distinctions, and Interconnections. In *Spatial Cognition*, Christian Freksa, Christopher Habel, and Karl F. Wender (Eds.). Lecture Notes in Artificial Intelligence, Vol. 1404. Springer, Berlin, 1–17.
- [22] Alexander Klippel, Carsten Dewey, Markus Knauff, Kai-Florian Richter, Daniel R. Montello, Christian Freksa, and E.-A. Loeliger. 2004. Direction Concepts in Wayfinding Assistance Systems. In *Workshop on Artificial Intelligence in Mobile Systems 2004 (AIMS 2004)*, Joerg Baus, Christian Kray, and Robert Porzel (Eds.), Vol. Memo Nr. 84. Sonderforschungsbereich 378, Saarbrücken, Germany, 1–8.
- [23] Alexander Klippel and Daniel R. Montello. 2007. Linguistic and Non-Linguistic Turn Direction Concepts. In *Spatial Information Theory*, Stephan Winter, Matt Duckham, Lars Kulik, and Benjamin Kuipers (Eds.). Lecture Notes in Computer Science, Vol. 4736. Springer, Berlin, 354–372.
- [24] Werner Kuhn. 2012. Core concepts of spatial information for transdisciplinary research. *International Journal of Geographical Information Science* 26, 12 (2012), 2267–2276.
- [25] Stephen C. Levinson. 1996. Frames of Reference and Molyneux’s Question: Crosslinguistic Evidence. In *Language and Space*, Paul Bloom, Mary A. Peterson, Lynn Nadel, and Merrill F. Garrett (Eds.). The MIT Press, Cambridge, MA, 109–169.
- [26] Gérard Ligozat. 2013. *Qualitative Spatial and Temporal Reasoning*. John Wiley & Sons, Inc., Hoboken, NJ. <https://doi.org/10.1002/9781118601457>
- [27] Fernando Melo and Bruno Martins. 2017. Automated Geocoding of Textual Documents: A Survey of Current Approaches. *Transactions in GIS* 21, 1 (2017), 3–38. <https://doi.org/10.1111/tgis.12212>
- [28] Daniel R. Montello. 1993. Scale and Multiple Psychologies of Space. In *Spatial Information Theory*, Andrew U. Frank and Irene Campari (Eds.). Lecture Notes in Computer Science, Vol. 716. Springer, Berlin, 312–321.
- [29] Daniel R. Montello, Michael F. Goodchild, Jonathan Gottsegen, and Peter Fohl. 2003. Where’s Downtown? Behavioral Methods for Determining Referents of Vague Spatial Queries. *Spatial Cognition and Computation* 3, 2&3 (2003), 185–204.
- [30] Reinhard Moratz and Thora Tenbrink. 2006. Spatial Reference in Linguistic Human-Robot Interaction: Iterative, Empirically Supported Development of a Model of Projective Relations. *Spatial Cognition & Computation* 6, 1 (2006), 63–107. https://doi.org/10.1207/s15427633scc0601_3
- [31] D. Punjani, K. Singh, A. Both, M. Koubarakis, I. Angelidis, K. Bereta, T. Beris, D. Bilidas, T. Ioannidis, N. Karalis, C. Lange, D. Pantazi, C. Papaloukas, and G. Stamoulis. 2018. Template-Based Question Answering over Linked Geospatial Data. In *Proceedings of the 12th Workshop on Geographic Information Retrieval (GIR’18)*. ACM, New York, NY, USA, Article 7, 10 pages. <https://doi.org/10.1145/3281354.3281362>
- [32] Ross Purves, Stephan Winter, and Werner Kuhn. 2019. Places in information science. *Journal of the Association for Information Science and Technology* 70, 11 (2019), 1173–1182. <https://doi.org/10.1002/asi.24194>
- [33] Ross S Purves, Paul Clough, Christopher B Jones, Mark H Hall, Vanessa Murdock, et al. 2018. Geographic information retrieval: Progress and challenges in spatial search of text. *Foundations and Trends in Information Retrieval* 12, 2-3 (2018), 164–318.
- [34] Marco Ragni and Markus Knauff. 2013. A theory and a computational model of spatial reasoning with preferred mental models. *Psychological Review* 120, 3 (2013), 561–588. <https://doi.org/10.1037/a0032460>
- [35] Afshin Rahimi, Trevor Cohn, and Timothy Baldwin. 2017. Continuous Representation of Location for Geolocation and Lexical Dialectology using Mixture Density Networks. In *Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Copenhagen, Denmark, 167–176. <https://doi.org/10.18653/v1/D17-1016>
- [36] Kai-Florian Richter, Stephan Winter, and Urs-Jakob Rüetschi. 2009. Constructing Hierarchical Representations of Indoor Spaces. In *Tenth International Conference in Mobile Data Management, Workshop on Indoor Spatial Awareness*, Yu-Chee Tseng, Peter Scheuermann, and Ralf Hartmut Güting (Eds.). IEEE Press, Taipei, Taiwan, 686–691. <https://doi.org/10.1109/MDM.2009.117>

- [37] Hanan Samet, Jagan Sankaranarayanan, Michael D. Lieberman, Marco D. Adelfio, Brendan C. Fruin, Jack M. Lotkowski, Daniele Panozzo, Jon Sperling, and Benjamin E. Teitler. 2014. Reading news with maps by exploiting spatial synonyms. *Commun. ACM* 57, 10 (2014), 64–77. <https://doi.org/10.1145/2629572>
- [38] Bill Schilit, D. M. Hilbert, and J. Trevor. 2002. Context-aware communication. *IEEE Wireless Communications* 9, 5 (2002), 46–54. <https://doi.org/10.1109/MWC.2002.1043853>
- [39] Albrecht Schmidt, Michael Beigl, and Hans-W. Gellersen. 1999. There is more to context than location. *Computers & Graphics* 23, 6 (1999), 893–901. [https://doi.org/10.1016/S0097-8493\(99\)00120-X](https://doi.org/10.1016/S0097-8493(99)00120-X)
- [40] Dan Sperber and Deirdre Wilson. 2004. Relevance Theory. In *Handbook of Pragmatics*, L. Horn and G. Ward (Eds.). Blackwell, Oxford, 607–632.
- [41] Thora Tenbrink and Werner Kuhn. 2011. A model of spatial reference frames in language. In *Spatial Information Theory*, Max Egenhofer, Nicholas A. Giudice, Reinhard Moratz, and Michael Worboys (Eds.). Springer, Berlin, 371–390.
- [42] Eran Toch, Yang Wang, and Lorrie Faith Cranor. 2012. Personalization and privacy: a survey of privacy risks and remedies in personalization-based systems. *User Modeling and User-Adapted Interaction* 22, 1 (2012), 203–220. <https://doi.org/10.1007/s11257-011-9110-z>
- [43] Maria Vasardani, Sabine Timpf, Stephan Winter, and Martin Tomko. 2013. From Descriptions to Depictions: A Conceptual Framework. In *Spatial Information Theory*, Thora Tenbrink, John Stell, Antony Galton, and Zena Wood (Eds.). Lecture Notes in Computer Science, Vol. 8116. Springer, Cham, 299–319.
- [44] Mark Weiser. 1991. The Computer for the Twenty-First Century. *Scientific American* 265, 3 (1991), 94–105.
- [45] Stephan Winter and Christian Freksa. 2012. Approaching the Notion of Place by Contrast. *Journal of Spatial Information Science* 2012, 5 (2012), 31–50.
- [46] Stephan Winter, Martin Tomko, Maria Vasardani, Kai-Florian Richter, Kourosh Khoshelham, and Mohsen Kalantari. 2019. Infrastructure-Independent Indoor Localization and Navigation. *Comput. Surveys* 52, 3 (2019), 61:1–61:24.
- [47] Shanyang Zhao. 2003. Toward a taxonomy of copresence. *Presence: Teleoperators & Virtual Environments* 12, 5 (2003), 445–455.