Spatial Interactions between Humans and Assistive Agents

Franz J. Kurfess and Greg Flanagan
California Polytechnic State University
San Luis Obispo, CA, U.S.A.

Mehul Bhatt[*]
SFB/TR 8 Spatial Cognition
University of Bremen, Germany

Abstract

While computers assist humans with tasks such as navigation that involve spatial aspects, agents that can interact in a meaningful way in this context are still in their infancy. One core issue is the mismatch in the representation of spatial information a computer-based system is likely to use, and the one a human is likely to use. Computers are better suited for quantitative schemes such as maps or diagrams that rely on measurable distances between entities. Humans frequently use higher-level, domain-specific conceptual representations such as buildings, rooms, or streets for orientation purposes. Combined with the person-centric world view that we often assume when we refer to spatial information, it is challenging for agents to convert statements using spatial references into assertions that match their own internal representation. In this paper, we discuss an approach that uses natural language processing and information extraction tool kits to identify entities and statements about their spatial relations. These extractions are then processed by a spatial reasoner to convert them from the human conceptual space into the quantitative space used by the computer-based agent.

Motivation and Background

This paper investigates the domain of human-agent interaction in spatial contexts. As an example, we use a scenario where a person explores an exhibit at a museum, supported by an agent that offers guidance about the layout and contents of the exhibit. Our emphasis here is on the spatial structure of the museum, and the arrangement of the pieces on display, not on information about the pieces themselves. We also make the assumption that human and agent communicate in a possibly restricted natural language in written or spoken form, or a mixture of both. For the sake of simplicity, we assume that the input to the agent is provided as text, so as to avoid the complications of speech recognition. The core issues of our scenario are

- communication about spatial aspects of the environment
- spatial aspects of co-existence in environments

- planning and coordination of activities with spatial constraints

In the simplest case, an instance of the agent could be a smart phone with sensors that convey visual information (e.g. a camera) and localization information (such as GPS or a WiFi-based system), combined with a map of the building and rooms that host the exhibit. In this case, the mobility of the agent is dependent on the human, and the operation of the camera also requires human assistance. A more sophisticated agent could be a robot that acts as a "tour guide" for the human. This agent has a significantly higher degree of autonomy, and also a more capable set of sensors. In either case, the agent can guide the human through the exhibit on a pre-defined or customized tour, but also respond to queries about the exhibit and the layout of the building (e.g. "Where are the bathrooms?", "What's the quickest way to the Mona Lisa?"). in the following, we examine the interaction between humans and agents in this scenario, with particular emphasis on spatial aspects. The work described here is under partial implementation in research projects on information extraction (Assal et al. 2010) and spatial cognition within assistance systems (Bhatt 2010b; Flanagan 2010; Bhatt and Guesgen 2009; Guesgen and Bhatt 2010; Bhatt and Freksa 2010).

The work discussed here combines aspects from two areas: Interactions between agents, in particular communication between computer-based assistive agents and humans, and dealing with spatial aspects of the environment in which the agents operate. In this section, we will examine related work in a few fields most relevant for our overall purpose.

Human-Agent Interaction

When humans interact with computer-based systems, the current interaction methods are typically constrained by the capabilities and limitations of computer systems, and the onus is on the human counterpart to adapt to those constraints. This is acceptable under many circumstances, especially if computers are used as tools to support specific tasks such as viewing or writing documents, or organizing information in a data base or spreadsheet. Considering the physical interaction space in which these exchanges take place, the human and computer counterparts are typically within a distance that allows humans to use their hands to interact, and to view items displayed on a screen com-

fortably. Most interactions are done in a stationary setup, where the locationz of computer and human are stable (e.g. centered around a desk). With laptops and mobile devices, a much higher degree of mobility is achieved, but the device still must be within reaching distance of the human for comfortable use. Using *autonomy* as a distinguishing feature between agents and conventional programs (Russell and Norvig 2009; Franklin and Graesser 1996), the close proximity for interaction between human and agent can not be maintained, and interaction must be enabled over somewhat larger distances. In our context, for example, we can assume that human and agent are within mutual visibility and hearability range (often in the same room), but not close enough for direct contact interaction. Better candidates for interaction methods under such circumstances are ones that rely on vision and hearing as sensory modalities since their range allows some distance between the agents involved. Many robots as physical instances of agents include sensors and actuators that in principle allow such communication, although in practice there may be significant obstacles and limitations. Often such robots are equipped with cameras that allow the use of vision as an essential sensory channel for navigation, obstacle avoidance, object recognition and manipulation. Vision also can be used for communication purposes, e.g. through gestures, sign language, or lip reading. Sign language requires special skills that few humans have, and reading sign language is far from trivial for computer-based systems. Lip reading allows humans to use spoken natural language, but it requires a direct line of sight, and is practical only in combination with auditory speech recognition. Gestures are quite natural for humans, and relate well to spatial aspects, e.g. by pointing at objects, or indicating directions. However, as a communication language, gestures have a limited expressiveness: There is a limited vocabulary, and beyond a small set of reasonably unambiguous ones gestures are not very well agreed upon. So in our context, the most practical communication method relies on spoken natural language, perhaps with a limited vocabulary and grammatical structure. It allows humans and agents to communicate as long as they are within hearing range, and from a human perspective it has the great advantage that it is very natural as a communication method. Obviously it is not without disadvantages: It is not practical in noisy environments, may not be appropriate for situations that require silence, and it imposes relatively heavy computational requirements on the computer-based system.

Natural Language Processing and Information Extraction

Considering that applications are available for smartphones that perform voice recognition acceptably well, our assumption is that the agent receives information from the user in a text-based format; we will assume that our computer-based agent has basic speech recognition capabilities, including the ability to recognize terms that deal with spatial aspects of the environment. Through Natural Language Processing and Information Extraction toolkits such as OpenNLP, the Stanford Named Entity Recognizer, and Thomson Reuters' OpenCalais (Apache 2010; OpenCalais 2010; Stanford NER 2010; OpenNLP 2010), essential information such as named entities (names of people, objects, buildings, etc.) and relations between these entities is extracted from a statement. For our particular emphasis on spatial aspects, the above toolkits may have to be enhanced to make sure that terms dealing with space and spatial relations will not be omitted. Since the terminology here is reasonably confined, this can be done through a combination of pattern matching via regular expressions and a thesaurus like Wordnet(WordNet 2010), or an ontology of spatial terms (Grenon and Smith 2004). Two important goals from our perspective are to establish a shared meaning for spatial terms between the agents (Williams 2004), and to identify entities of interest in the environment, such as features of the room (doors, windows), obstacles, other agents (human or computer-based), or objects relevant to the task at hand.

Spatial Assistance Systems

The broad perspective of this paper is rooted in our interpretation of a general class of assistance systems concerned with the assistance, assurance and empowerment of humans in their everyday professional and personal lives (Bhatt, Schultz, and Freksa 2010). Confining the discussion to the spatial domain, these are systems of human-computer interaction involving the representation of space from several different perspectives — the psycholinguistic and cognitive perspective of humans, the abstract knowledge-centric (symbolic) perspective of an intelligent software construct, and the all too complex and inherently numerical or quantitative perspective of the computer (read: real world) (Bhatt, Schultz, and Freksa 2010; Bhatt and Freksa 2010). In our terminology, spatial assistance systems are cognitive agents that 'know' the properties of physical space and are skilled to deal with them in such a way that they can support humans. These might need help, either because they are less skilled, or because they want to solve spatial problems that require more than one agent, or simply, because they want to pursue other activities, instead. A special requirement for spatial assistance system is that they are able to empathise with their human partners to a certain extent; i.e. they should adapt to the needs of people rather than require people to adapt to their needs. This manner of assistance inherently involves fundamental representational and computational challenges, which also inspires the broad research questions underlying approach for *reasoning about space, actions and change* (Bhatt 2010b; Bhatt, Guesgen, and Hazarika 2010).

Spatial Cognition within Assistance Systems

Extracting named entities and spatial terms may be sufficient for an agent to answer simple queries like "Where is `object X`?", where `object X` is an item on display, listed in the exhibition catalog, and thus a known entity for the agent. It is not sufficient, however, for other queries with spatial information, such as "What is the tall statue to the right of the door?" This requires qualitative spatial reasoning, including consideration of the location and orientation of the speaker, the relative position of one entity (the statue) with respect to another (the door),
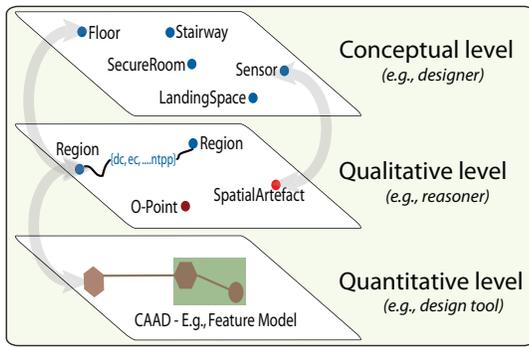
(a) Key (Ontological) Perspectives
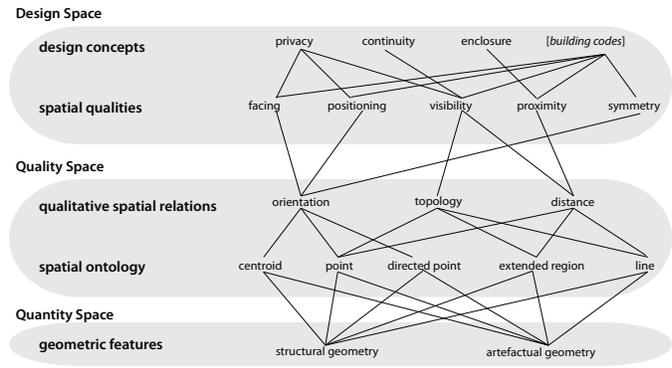
(b) Multiple Modalities

Figure 1: Multi-Perspective Representational Semantics. E.g., the spatial design domain (Bhatt and Freksa 2010)

and issues like visibility (Bhatt, Dylla, and Hois 2009; Flanagan 2010).

One of the critical issues here is the mapping between the conceptual model of the human, and the respective world model of the agent. In the agent's case, the underlying representation scheme of its world model is likely to be a map or an architectural plan of the building, augmented by information about the objects in the exhibit (such as their locations). While a human visitor may also have a "map" conceptual model of the museum, it may require significant mental effort to translate a statement into this framework of reference, and may only be realistically accurate with the aid of an actual physical map. So the above statement would have to be translated into something like "What is the statue in Room 25, on the eastern side of the door that leads into Room 18?" Clearly this is not only inconvenient for the user, but also increases the complexity of the natural language statement to by analyzed by the agent. Using tools like the *Conceptual Requirements Reasoner* (Flanagan 2010), concepts from the design space (the perspective of the user) are converted into the quality space (qualitative spatial relations and spatial abstraction terms), and finally into the quantity space used by the agent (structural and artefactual geometry as represented in the map or architectural plan). Figure 1 illustrates the concept of multi-perspective representational semantics for the domain of spatial design (Bhatt and Freksa 2010):

Conceptual level (design space). The *design space* addresses domain aspects at the conceptual level (in this case architecture and museum design) such as enclosure, continuity, and privacy, and high-level spatial qualities such as facing, positioning, visibility, and proximity. These concepts and qualities form the basis for a qualitative based descriptor language that describes various experiential aspects of the architectural design and can be used for reasoning about security systems, evacuation routes or guided tours.

Qualitative level (spatial ontology). The *quality space* level includes qualitative spatial relations (orientation, topology, and distance) and spatial abstractions (point, directed point, line, convex hull) that are used as the basis for building the constructs at the conceptual design level and play an intermediating role between the conceptual level of a design

and its precise quantitative floor plan. Figure 1(b) shows a multi-hierarchical structure that maps the concepts in the design space to the geometric representations in the quantity space. The Conceptual Requirements reasoner incorporates logical rules from Qualitative Spatial Reasoning frameworks such as the Single-Cross Calculus (Freksa 1992) for intrinsic orientational reasoning, Oriented Point Relation Algebra (Moratz 2006) for extrinsic orientational reasoning, and the Region Connection Calculus (Randell, Cui, and Cohn 1992) for topological reasoning and allows the representation of and reasoning about space in a formal, qualitative framework. The crux of using such a framework is that it does not come with the overhead of specific quantitative representations used in conventional maps and design diagrams that incorporate actual measurements. Obviously, for some purposes (such as calculating distances) the quantitative representations are more appropriate and even necessary.

Quantitative level (the real world). While it is quite challenging for such an agent to convert a simple statement uttered by a human user, in this and related scenarios it is usually not too difficult to provide a response to such a statement in spatial terms. Frequently, it will be sufficient to display the information requested on a map or schematic diagram, and the user can translate this into the real environment with moderate effort. Even if the response is to be formulated in text, the hard work has already been done: The relative positions of the human user, the objects addressed in the request, and the reference objects on the map have already been identified. This can then be converted into a series of statements that give the user directions to the requested goal, for example. It also can serve as the basis for actions by the agent, such as moving in a particular direction or to a specific location, or performing an action involving a referenced object.

This also completes the communication and interaction cycle: Starting with the communication by exchanging signals such as spoken text in the physical world and observing aspects of the real world, the relevant spatial relations and abstractions are identified by the agent at the qualitative level, and a connection is established with domain aspects at the conceptual level. The response of the agent then again may refer to qualitative aspects such as orientation or

distance of entities, and their actual placement in the real world.

Framework and Implementation

Within the SFB/TR 8 Spatial Cognition effort, a consortium of universities and research institutions in Germany has been developing a computational framework for spatial reasoning within a wide-range of *Spatial Assistance Systems* (Bhatt, Schultz, and Freksa 2010; Bhatt 2010a). In this context, one of the initial methodological techniques has been the use of spatial reasoning within a Constraint Logic Programming framework to satisfy design requirements for a rather specific domain, namely the design of art museums (Flanagan 2010). This approach relies on a *Conceptual Requirements Reasoner* to validate design requirements formulated by architects and museum designers. The main purpose is to show the process of representing and reasoning about design requirements in a formal computational framework.

At a high level, a museum designer may define the following requirements:

- Maximize visitor utilization of exhibitions

- Encourage free-flowing exploration throughout the museum

- Adhere to museum requirements of accessibility and security

The first one can be broken down into factors that influence the movement patterns of visitors, such as the positioning of doorways, spatial arrangement of display cases, positioning of furniture and statues, and congestion. A further refinement leads to a requirement stating that exhibition doorways should be positioned on opposing sides of gallery rooms. At this level, we can relate entities in the requirement (doorways, rooms) through spatial relations (opposing sides).

The design hierarchy used in the reasoner corresponds to the three levels shown in Figure 1(b). At the conceptual level, architectural concepts are built using qualitative spatial attributes (QSA) found in architecture. Spatial relationships of orientation, topology, and distance can be used, as well as other geometric primitives such as area, angle, and length measurements. These qualitative spatial relationships emerge from the quantitative, physical, and artefactual geometries of the the design, representing the physical elements of the building (doors, walls, columns, windows, etc.), artefactual extensions (functional space, operational space, range space, etc.) and interior design elements (furniture, lights, decorations, etc.).

QSAs can be perceptual qualities that are described from a specific vantage point, or intrinsic qualities that are inherent in the buildings spatial structure. Theones considered here are positioning, facing, visibility, proximity, and symmetry.

The facing attribute can be split up into two cases, facing-towards and facing-away. Facing-towards indicates that the directed point A is oriented towards another point B. Facing-away is the opposite, and can be expressed simply as the negation of Facing-towards. Using an Oriented Point Relation Algebra (OPRA) (Moratz 2006), the space surrounding a point is divided into sectors, illustrated by lines emanating from the points in Figure 2. Then the fact that point A
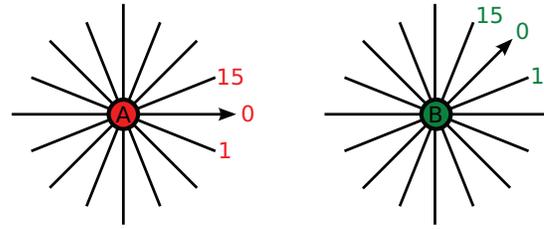

Figure 2: "Facing" Attribute

is facing-towards point B can be expressed by specifying in which sector of A's space B can be found. For a space divided into n sectors, we can specify that A faces towards B if B is in sectors $\{-k, \ldots, 0, \ldots, +k\}$ of A (where k is a small number), as depicted in the figure; B, conversely, is facing away from A since A is not within those sectors of B. These two cases are captured in the following Prolog code snippet.

```
facing_towards(Obj1, Obj2) :-
        (opra(Obj1, Obj2, 0) ;
        opra(Obj1, Obj2, 1) ;
        opra(Obj1, Obj2, 31)).
facing_away(Obj1, Obj2) :-
        not facing_towards(Obj1, Obj2).
```

The positioning attribute specifies the orientation of a point C (referent) with respect to a point B (relatum) from the viewpoint of a point A (origin). Then from A's perspective, C can be described as in front of (same side) or behind B (opposing side), or to the left or right of B. In this case, the Single Cross Calculus (Freksa 1992) serves as the basis of the relative positions of the points. The positioning relationships are then also represented as Prolog rules, with a separate rule for each relationship: opposing side, same-side, left-side, and right-side.

In a similar manner, the other attributes are broken down into rules that specify the relationships between the objects under consideration. Then a requirement like the one mentioned above referring to opposing sides can be validated through these rules. These examples illustrate how spatial relations used by humans at the conceptual level can be expressed as relationships between sets of points at the qualitativel level, which can be implemented in a reasonably straightforward manner as spatial reasoners in Prolog and similar languages. To use them in realistic environments also requires a quantitative description of the environment, such as an architectural drawing, a floor plan, or a map constructed by a robot via exploration. These qualitative descriptions may require some adaptation for use by the reasoner, which can be done either through the tools that are used to generate and store such descriptions (e.g. CAD tools), or through preprocessing scripts (Flanagan 2010).

Spatial Considerations in Museum Design

At this point, the case study focuses on the examination and validation of design requirements by relating high-level concepts from the application domain to qualitative descriptions of the physical environment. This does not involve meaningful interaction between humans and computer-based agents in order to assist humans. The system implemented, how-

ever, can then be incorporated into a museum guide application for a smartphone, for example, that is capable of dealing with spatial expressions in queries from its user. In this case, the high-level requirements expressing necessary or desirable features of a museum are replaced by statements reflecting the experience, needs, and desires of the user. So a stattement like "I can't see the sculpture of the Degas dancer. Is it in this room, or the one behind it?" requires the consideration of enclosure ("in this room"), visibility ("can't see"), positioning ("behind"), location (where is the visitor), orientation (which direction is the visitor facing), and knowledge about specific artefacts ("the Degas dancer") and their location. As a shortcut, it may be sufficient for a museum guide app to perform named entitiy extraction ("Degas dancer sculpture"), look it up in the museum catalog, and display it on a map. While this reduces the need for spatial reasoning, it puts the onus of relating the entitites on the map to their real-world counterparts on the human user.

Validation and Evaluation

Our emphasis currently lies on demonstrating the feasibility of using this approach to bridge the gap between the conceptual-level thinking and reasoning used by humans, and the "grounding" needed to establish an agent in a real environment. Thus, the main validation will be achieved by running scenarios in the museum environment to the spatial reasoner, and check the results for correctness and plausibility. At the moment, the main validation method is to inspect the output generated by the reasoner. An extension of this manual check will be the use of multiple reasoners, possibly based on different spatial calculi, where the output generated by one reasoner is examined by another reasoner. This will be augmented by visualizing the objects and relationships under investigation in a simulated environment, probably by incorporating additional layers in a map or floorplan. At the same time, we are examining candidates for metrics to measure the quality of the translations between the three levels of abstraction. Our medium-term plan is to convert the current system from a conventional computer environment (desktop or laptop) into one that can be installed in a robotic agent, or an agent in a virtual environment such as Second Life or a game environment such as Microsoft's Xbox witht he Kinect sensor.

Further Application Domains

Spatial information plays an important role in many aspects of our lives, be it the arrangement of objects in our home, navigating a familiar or unfamiliar terrain, or playing games in virtual environments.

Geographic Information Systems often rely on quantitative representations such as maps. One of the major advantages of such a representation is the preservation of relative distances, allowing humans to quickly estimate distances or travel times. In some situations, however, qualitative information is more essential, and requires careful consideration of the interaction between humans and computers. Car navigation systems (and smart phones incorporating similar functionality) are a good example for this, and their trans-

lation between quantitative ("Turn left after 300 feet") and qualitative statements ("Turn left at the next traffic light") can be confusing (Schultz, Guesgen, and Amor 2006).

Vehicle traffic will include more and more vehicles that have a certain degree of autonomy, ranging from distance-sensitive cruise control over lane following to the fully autonomous vehicles as recently demonstrated by Google. The interaction between the agent (the vehicle) and the user must bridge similar gaps between the conceptual space of the user (their mental map of a city or region) and the quantitative representation the vehicle mostly relies on (Terziyan, Kaykova, and Zhovtobryukh 2010; Hsu et al. 2010; Lertlakkhanakul, Hwang, and Choi 2009; Vales-Alonso et al. 2008). Humans often rely on visual landmarks for orientation, and tend to combine them with qualitative spatial references in statements like "Turn sharp right at the intersection right after the church with the high steeple on the right". In addition to the linguistic challenge of disambiguating between the three instances of "right" in this sentence, an agent will need information beyond a typical map to be able to follow such an instruction.

Smart homes also will see agents that need to be able to convert user statements about spatial aspects of tasks into their own representation, which likely will utilize a floor plan or map of the home (Uhm et al. 2010), (Brdiczka, Crowley, and Reignier 2009), (Lertlakkhanakul, Choi, and Kim 2008), (Bhatt and Guesgen 2009; Gottfried, Guesgen, and Hübner 2006), (Augusto and Nugent 2006). Wouldn't it be nice to tell your Roomba "Hey, be carefull when you vacuum under the desk in the corner by the window! Lots of cables under that desk ..."

One of the earliest domains to deal with spatial aspects of human-agent interaction may well be entertainment (Schlieder, Kiefer, and Matyas 2006). Computers as well as game consoles not only provide stimulating virtual environments with some realistic spatial aspects, they also incorporate rich interfaces such as Microsofts Xbox with the Kinect module (Microsoft 2010). The user is immersed into the virtual environment, and uses the agent's sensors (e.g. keyboard, mouse, controller, camera, microphone) to navigate. From the agent's perspective, this is very convenient since the user takes care of most of the translation from the user's design space (mental model) into the agent's representation of the virtual world. While there are spatial aspects to consider if natural language is used for the interaction ("Watch out, there's a ghost on your left!"), the agent does not have to deal with the physical environment in which the user is situated. This may change with the rich interfaces that can obtain a realistic view of the user's environment (e.g. their living room). Such devices also have the computational power for decent voice recognition, and may enable users to express their interactions in spoken natural language.

So one of these days I'll hopefully be able to interact with my personal agent: "Hi Robbie, please bring me the cold beer from the back of the fridge." Robbie pulls a beer bottle out of the fridge, and brings it over to the couch. "Robbie, how long have you been working for me? You should know by now that I don't like Becks; please bring me my Optimator instead!"

References

Apache. 2010. Apache UIMA. http://uima.apache.org.

Assal, H.; Seng, J.; Kurfess, F.; Schwarz, E.; and Pohl, K. 2010. Partnering enhanced-nlp with semantic analysis in support of information extraction. In *2nd International Workshop on Ontology-Driven Software Engineering (ODiSE 2010) at the ACM SPLASH 2010 Conference*.

Augusto, J. C., and Nugent, C. D., eds. 2006. *Designing Smart Homes, The Role of Artificial Intelligence*, volume 4008 of *Lecture Notes in Computer Science*. Springer.

Bhatt, M., and Freksa, C. 2010. Spatial computing for design: An artificial intelligence perspective. In *NSF International Workshop on Studying Visual and Spatial Reasoning for Design Creativity (SDC'10)*. http://cindy.informatik.uni-bremen.de/cosy/staff/bhatt/seer/Bhatt-Freks%a-SDC-10.pdf.

Bhatt, M., and Guesgen, H., eds. 2009. *Spatial and Temporal Reasoning for Ambient Intelligence Systems (STAMI 09)*. SFB/TR 8 Spatial Cognition Report Series, No. 020-08/2009.

Bhatt, M.; Dylla, F.; and Hois, J. 2009. Spatio-terminological inference for the design of ambient environments. In Hornsby, K.; Claramunt, C.; Denis, M.; and Ligozat, G., eds., *Spatial Information Theory*, volume 5756 of *Lecture Notes in Computer Science*. Springer Berlin / Heidelberg. 371–391.

Bhatt, M.; Guesgen, H.; and Hazarika, S., eds. 2010. *Spatio-Temporal Dynamics (STeDy 10)*. ECAI Workshop Proceedings., and SFB/TR 8 Spatial Cognition Report Series. http://cindy.informatik.uni-bremen.de/cosy/events/ecai10/.

Bhatt, M.; Schultz, C.; and Freksa, C. 2010. The 'Space' in Spatial Assistance Systems: Conception, Formalisation and Computation. (in review: available upon editorial / reviewer request).

Bhatt, M. 2010a. Commonsense inference in dynamic spatial systems: Epistemological requirements. In Guesgen, H. W., and Murray, R. C., eds., *FLAIRS Conference*. AAAI Press.

Bhatt, M. 2010b. Reasoning about space, actions and change: A paradigm for applications of spatial reasoning. In *Qualitative Spatial Representation and Reasoning: Trends and Future Directions*. IGI Global, USA. http://cindy.informatik.uni-bremen.de/cosy/staff/bhatt/seer/Bhatt-2010-%RSAC-Book.pdf.

Brdiczka, O.; Crowley, J.; and Reignier, P. 2009. Learning situation models in a smart home. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on* 39(1):56 –63.

Flanagan, G. 2010. Spatial cognition and natural language (working title). Master's thesis, California Polytechnic State University, San Luis Obispo, CA, U.S.A.

Franklin, S., and Graesser, A. 1996. Is it an agent, or just a program?: A taxonomy for autonomous agents. 21–35. Springer-Verlag.

Freksa, C. 1992. Using orientation information for qualitative spatial reasoning. In Frank, A.; Campari, I.; and Formentini, U., eds., *Theories and Methods of Spatio- Temporal Reasoning in Geographic Space*, volume 639 of *LNCS*, 162–178. Berlin: Springer-Verlag.

Gottfried, B.; Guesgen, H. W.; and Hübner, S. 2006. Spatiotemporal reasoning for smart homes. In Augusto and Nugent (2006), 16–34.

Grenon, P., and Smith, B. 2004. Snap and span: Towards dynamic spatial ontology. *Spatial Cognition & Computation: An Interdisciplinary Journal* 4(1):69–104.

Guesgen, H. W., and Bhatt, M. 2010. Advances in spatial and temporal reasoning. *J. UCS* 16(11):1388–1389.

Hsu, I. Y.-Y.; Wodczak, M.; White, R. G.; Zhang, T.; and Russell Hsing, T. 2010. Challenges, approaches, and solutions in intelligent transportation systems. In *Second International Conference on Ubiquitous and Future Networks (ICUFN), 2010*, 366 –371.

Lertlakkhanakul, J.; Choi, J. W.; and Kim, M. Y. 2008. Building data model and simulation platform for spatial interaction management in smart home. *Automation in Construction* 17(8):948 – 957.

Lertlakkhanakul, J.; Hwang, S.; and Choi, J. 2009. Developing spatial information framework for urban computing environment: A socio-spatial computing framework for smart urban environment. In *International Conference on Management and Service Science (MASS '09)*, 1 –5.

Microsoft. 2010. Kinect. www.xbox.com/kinect.

Moratz, R. 2006. Representing relative direction as a binary relation of oriented points. In *Proceeding of the 2006 conference on ECAI 2006*, 407–411. Amsterdam, The Netherlands, The Netherlands: IOS Press.

OpenCalais. 2010. Opencalais. http://www.opencalais.com.

OpenNLP. 2010. OpenNLP. http://opennlp.sourceforge.net/.

Randell, D. A.; Cui, Z.; and Cohn, A. G. 1992. A spatial logic based on regions and connection. In *PROCEEDINGS 3RD INTERNATIONAL CONFERENCE ON KNOWLEDGE REPRESENTATION AND REASONING*.

Russell, S. J., and Norvig, P. 2009. *Artificial Intelligence: A Modern Approach*. Upper Saddle River, NJ: Prentice Hall, 3. edition.

Schlieder, C.; Kiefer, P.; and Matyas, S. 2006. Geogames: Designing location-based games from classic board games. *Intelligent Systems, IEEE* 21(5):40 –46.

Schultz, C. P. L.; Guesgen, H. W.; and Amor, R. 2006. Computer-human interaction issues when integrating qualitative spatial reasoning into geographic information systems. In Billinghurst, M., ed., *CHINZ*, volume 158 of *ACM International Conference Proceeding Series*, 43–51. ACM.

Stanford NER. 2010. Stanford named entity recognizer. http://nlp.stanford.edu/software/CRF-NER.shtml.

Terziyan, V.; Kaykova, O.; and Zhovtobryukh, D. 2010. Ubiroad: Semantic middleware for context-aware smart road environments. In *Fifth International Conference on Internet and Web Applications and Services (ICIW), 2010*, 295 –302.

Uhm, Y.; Lee, M.; Kim, Y.; Byun, J.; Kim, G.; Lee, B.; and Park, S. 2010. Development of portable intelligent gateway system for ubiquitous entertainment and context-aware services. In *Digest of Technical Papers International Conference on Consumer Electronics (ICCE), 2010*, 477 –478.

Vales-Alonso, J.; Egea-Lopez, E.; Muoz-Gea, J.; Garcia-Haro, J.; Belzunce-Arcos, F.; Esparza-Garcia, M.; Perez-Maogil, J.; Martinez-Alvarez, R.; Gil-Castieira, F.; and Gonzalez-Castao, F. 2008. Ucare: Context-aware services for disabled users in urban environments. In *The Second International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies (UBICOMM '08)*, 197 –205.

Williams, A. B. 2004. Learning to share meaning in a multi-agent system. *Autonomous Agents and Multi-Agent Systems* 8(2):165–193.

WordNet. 2010. Wordnet. http://wordnet.princeton.edu.