Telling Rolland where to go: HRI dialogues on route navigation

Hui Shi and Thora Tenbrink[1]

SFB/TR8 Spatial Cognition, University of Bremen, Germany

Abstract

Our aim is to enable uninformed users to instruct the Bremen autonomous wheelchair Rolland to move towards specific goals via route directions. Empirical data are used to augment the robot's internal representation as well as to identify a number of conceptual problem areas involved in the spatial communication between a human and a robot. The problematic aspects are tackled via the design of a dialogue system capable of managing the available information and asking appropriate clarification questions in order to ensure efficient communication.

1. Introduction

Within the SFB/TR 8 on Spatial Cognition, one prominent research aim is to achieve effective and natural communication between humans and robots about spatial issues. While some of the work so far has addressed spatial instruction scenarios in which a robot was instructed to move towards one of several similar objects present in a scenario (e.g., Moratz & Tenbrink, in press; Tenbrink & Moratz 2003), the present aim is to enable previously uninformed users to "tell Rolland where to go" via a suitable dialogue model. 'Rolland' is the name of the Bremen autonomous wheelchair (e.g., Lankenau & Röfer 2000), a mobile robot equipped with sensorial and conceptual functionalities, which we use to investigate how humans communicate linguistically about route navigation. Comparable work of other research groups also dealing with the automatic treatment of route instructions given to robots is presented, for example, in Bugmann et al. (2004), Ligozat (2000), and Tschander

[1] A number of further people within the projects of the SFB/TR 8 on Spatial Cognition are involved in this work, for instance, John Bateman, Scott Farrar, Kerstin Fischer, Udo Frese, Bernd Krieg-Brückner, Christian Mandel, Reinhard Moratz, Thomas Röfer, Robert Ross, Tilman Vierhuff, and others. Their participation in the development of the ideas and results presented here is acknowledged.

et al. (2003). The specific features of route communication in a *dialogic* situation have seldom been investigated; however, a very thorough and fine-grained analysis of a Maptask dialogue is presented by Filipi and Wales (2004), who show that shifts of perspective (i.e., the conceptualization and linguistic representation of the spatial situation) occur systematically as required by the demands of the interaction. Such regular variability in speakers' representations needs to be accounted for in human-robot interaction, on the one hand by implementing a suitable "cognitive map" (cf. Section 3), and on the other hand by employing an adequate dialogue model that is capable of handling conceptual mismatches similar to the alignment and repair strategies in a natural dialogue.

The scenario we have in mind for the present work is as follows. Users are seated in the wheelchair and asked to move around in the university building, explaining the most important rooms and places (landmarks) to Rolland. They are told that the robot completes its internal map in this way. After that, the users are asked to instruct Rolland to navigate to a specific place where they have previously been together with the wheelchair. Rolland is then supposed to move autonomously to the intended goal. At this point, a number of communication problems may occur because of mismatches between the robot's knowledge and the user's linguistic representation. Thus, on the one hand we are interested in users' spontaneous strategies in instructing the robot for the navigation task; on the other hand we investigate how clarification dialogues can be initiated by the robot in an effective way in cases of communication failure.

In previous research, we already used rudimentary clarification dialogues, the design of which was based on the modules that failed to understand the robot's request (Moratz & Tenbrink, in press): If the lexicon failed, the robot told the user which word it could not understand. If the syntactic processing module failed, the robot asked the user to reformulate the instruction. If the computational model failed because the user employed a kind of spatial reference that the robot was not equipped to understand, the user was asked to choose a different spatial description. This simple dialogic model worked fairly well within a scenario consisting of only few objects situated within the range of vision. We take this as a starting point, but consider a route navigation scenario to be substantially more complex, therefore requiring a more sophisticated dialogue model. Furthermore, this rudimentary dialogue model was incapable of preventing users from getting frustrated with the robot because they did not succeed in finding out what the robot could understand. This happened in considerably fewer cases than in a previous experimental round in which there had not been any feedback at all except for "I don't understand", but naturally our aim is to ensure even smoother and more effective communication.

From our empirical data collected in a scenario in which the robot wheelchair did not act autonomously, we extract information provided by the participants in order to augment the robot's internal map, and we identify a number of conceptually problematic aspects that need to be dealt with in a real human-robot interaction scenario. To solve the problems, clarification dialogues are needed, for which purpose we suggest a suitable dialogue model that is flexible enough to capture the present insights as well as to be extended on the basis of future research.

2. Spatial reference in route navigation: An empirical study

The design of the present study is based on a method developed by K. Fischer (Fischer 2003), who also took part in the preparation and realisation of this study. One decisive feature of this methodology is to keep users basically uninformed about the robot's capabilities, and to suggest a realistic (though partly illusory) human-robot interaction scenario to them that is suitable for triggering the kind of utterances that the system should understand in the future. This method ensures that speakers' spontaneous linguistic strategies and choices can be investigated.

The 23 German and 6 English native speakers participating in our study were told that the robot wheelchair was capable of processing information given verbally in their native language, and that it would then use it for later navigation tasks. Rolland never gave any kind of response in reaction to the users' utterances, or feedback that it had actually understood the message. This is the case because in this exploratory phase Rolland's actual functionalities were not used at all, contrary to what we told the users.

The users were seated in the wheelchair and asked to steer Rolland around the hallways in the university along a prescribed route and talk to it, explaining the spatial relationships that Rolland would need to know in later tasks. In this phase, speakers displayed different individual priorities and strategies, and they produced descriptions on various levels of granularity. For example, they named entities along the route and described their spatial location, such as "the first door to the right here is the copy machine". In other cases, they described their current direction of movement, for instance (numbers indicate pauses in seconds): "left (5) straight ahead (20) and slightly left (3) slightly right (17) hard left", and the like. The results of this experimental phase were used later to augment the robot's internal map (cf. section 3).

In the next phase of the experiment, the users were asked to instruct Rolland to move to a particular place called the "Stugaraum" or "Stugeneck". We told them that Rolland would now be able to navigate autonomously to this place, which of course was not the case, since the robot could not integrate the information directly. In this phase, users produced in-advance route directions towards the goal, again using various strategies and employing different levels of granularity, but generally attending to the kinds of information they had previously given the robot. The results of this second phase were used as a data pool to identify potentially problematic conceptual phenomena (cf. section 4) which will need to be addressed in order for the human-robot interaction to succeed. For this purpose, we present a suitable dialogue model in section 5.

3. Cognitive map: a representation of Rolland's spatial knowledge

Humans' cognitive models of spatial situations do not directly mirror the actual relations. This is reflected in the ways humans represent their spatial concepts externally, for example, in verbal route directions or sketch maps (Tversky & Lee 1998). Such externalisations highlight how mental representations are schematised, systematically simplified, and often distorted (Tversky 2003, Talmy 2000); also, they reflect a range of different conceptualisations or "perspectives" (Tversky 1999, Filipi & Wales 2004). A number of investigations (e.g., Denis 1997, Michon & Denis 2001) classify the essential ingredients of a spatial description: they contain information about orientation and direction, about decision points, landmarks, and continuations of movements. Our own empirical data basically conform with this classification.

Within Artificial Intelligence, such mental conceptualisations of space are formalised and modeled in the subfield of qualitative spatial representation and reasoning. This kind of automation can make valuable predictions about human spatial behavior even in cases where a precise quantitative representation is not available or is too complicated to be treated computationally (Cohn et al. 1997). Spatial knowledge can be represented in a number of ways; in our approach we combine Freksa's Qualitative Spatial Calculus using orientation information (Freksa 1992) and the Route Graph (Werner et al. 2000) to represent Rolland's conceptual spatial knowledge for the communication with users.

Route Graph. The general concept of Route Graphs is suitable for modeling navigation of various kinds of agents in diverse scenarios (Krieg-Brückner et al. 2005). They can be used as metrical maps combining sensory input with metric computational models for controlling the navigation of robotic agents in the environment; likewise, they can be used at the cognitive level to abstractly model humans' topological knowledge while they act in the space. Importantly, Route Graphs integrate different routes as well as different kinds of information into a coherent graph-like structure.

The starting point for the definition is basic graph theory, yielding nodes and edges. A **node** of a Route Graph, called a *Place*, has a particular position and specification ("reference system") which may be related to a global reference system. An **edge** of a Route Graph is directed from a source node to a target node and is called a *route segment*. It has three additional attributes: an *entry*, a *course*, and an *exit.* The information associated with these attributes is specifically defined in each Route Graph instantiation. For example, in a route graph at the metrical level, an entry or an exit can be an angle defined by its degree value with respect to a global 2-D geometric system; the course is simply characterized by metrical data about length and width. On the other hand, an entry or exit at the cognitive level may contain qualitative orientation information (such as *to the left/right*), while the course consists of the path between two reorientations. Finally, a *Route* is defined by the conjunction of route segments from one place to another. The details of the general concept of Route Graphs are given in Werner et al. (2000) and Krieg-Brückner et al. (2005).

Qualitative spatial calculus using orientation information. Freksa's calculus for qualitative spatial reasoning (Freksa 1992) uses orientation information given by two points in 2-dimensional space, the start point and the end point of a movement. By combining the *front/back* and the *left/right* dichotomy, it distinguishes eight meaningful disjoint orientation relations (e.g., *straight-front, right-front*, *right-neutral* etc.). Freksa (1992) gives the following simple example to show how a location can be determined on the basis of our own location and other locations we know, which is an important application of orientation-based qualitative spatial reasoning.

Imagine walking from location *a* to location *c* and reaching the intermediate location *b*. We can describe orientation and distance of location *c* qualitatively with respect to the segment between location *a* and *b* as an oriented line, denoted as *vector* **ab**, i.e., we compare the vector **bc** to **ab** with respect to their orientation. At position *c*, we can compare the next vector **cd** with the previous one, **bc**. The inference step then determines the goal location *d* in relation to the initial vector **ab**.

Combining these two approaches we gain what we call the *conceptual route graph.* The conceptual route graph consists of the four basic types *orientation*, *landmark*, *place* and *vector.* Moreover, a set of relations between these types is defined. For example, *at* defines the relation between place and route mark, *on* the relation between place and vector, and *ori* the relation between two vectors or between a vector and a place. In a conceptual route graph, entries and exits of route segments are given as orientation relations, while courses of route segments are vectors. Fig.1 gives a conceptual route graph of the floor where our empirical study took place; the background shows the floor plan.

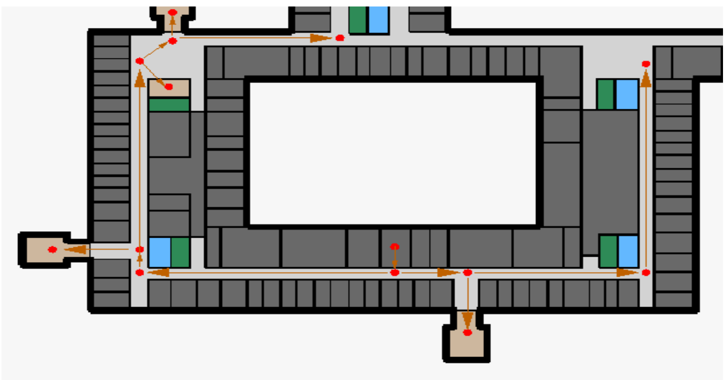

Fig. 1: A Conceptual Route Graph

There are several reasons for using the conceptual route graph. First, Freksa's qualitative spatial calculus is well-defined and provides an efficient way of representing and reasoning about causal actions and topological relations, while the Route Graph provides concepts like route segments, routes, and a number of spatial relations. Second, our empirical studies show that the combination is powerful enough to cover almost all of the users' descriptions used in the communication with the robot during the spatial navigation task. Third, in practice this representation supports an efficient mapping to the robot's spatial representation at the metric level for real navigation, and vice versa.

The results of the first experimental phase were added into the conceptual route graph. For example, Rolland "learned" the location of the mailbox room by the description "hier rechts von mir ist die Tür zu dem Raum mit den Postfächern" (here to the right of me is the door to the room with the mailboxes), represented as *{ori(v, p, right-neutral), at(theMailbox, p)}*, where *v* is the vector including the current place and orientation, *p* the place to the right of *v*, and the route mark "the mailbox" is at *p*.

4. Representation disparities

Our empirical data collected in the second experimental phase contain a number of situations in which the users' in-advance route descriptions cannot be understood or interpreted by Rolland according to its current knowledge representation, even after augmenting the internal representation. This section describes some of these problems.

Conceptual knowledge and reference resolution. Among the instructions we collected, a considerable number of utterances contains reference to entities, such as doors, rooms, or halls. An example is: "fahren an dem Kopierraum und dem Fächerraum vorbei" (drive past the copier room and the mailbox room). Here, the identifiability of the intended entities is presupposed, i.e., *at*-relations of the copier and the mailbox to a place should be defined, either because it was previously implemented or because it was added based on the user's utterances in the first phase.

With respect to vocabulary, implementation prior to experimentation is problematic as the actual kinds of references speakers make in a human-robot interaction scenario are hard to predict: users typically do not have access to either the lexicon, the internal representation, or to the perceptual abilities of the robot that they are currently dealing with, unless they are provided with a list of possible commands. Previous research has shown that speakers' assumptions concerning what robots can or cannot understand crucially determine their linguistic strategies towards the robot (Fischer & Moratz 2001). In our scenario, users could assume that the robot had been informed about the position of the entities by the prior exploration; however, it is not the case that they reliably used the kind of vocabulary that they had used in the exploration phase, although they quite obviously tried to do so (in one case, a speaker even tried to imitate her own previous pronounciation: “Rolland I want you to go to the stugaroom, or did I say sh-tugaroom”). For instance, the term “Fächerraum” (mailbox room) had not been previously used.

With respect to conceptual knowledge, the robot may not be capable of integrating some parts of the information. For example, the description “…vorbei an der Verwaltung von Herrn Müller…” (pass by the administration of Mr. Müller) cannot be straightforwardly resolved into “pass by the office of Mr. Müller of the administration”. Moreover, the user may refer to entities or actions in the real world that are present at some place in the robot's internal representation but not at the place the user refers to, which leads to a representation mismatch.

Underspecification. Many utterances contain vague and underspecified representations, such as “n kleines Stück geradeaus” (a bit straight on) instead of specifying an exact distance that the robot should cover, or a destination towards which it should move. Some users simply repeat direction instructions, e.g., “geradeaus geradeaus geradeaus (…) irgendwann nach links” (straight on straight on straight on (…) eventually to the left). Sometimes users employ expressions like ”ungefähr” (roughly), “weiß nicht” (don’t know), “irgendwo” (somewhere), reflecting uncertainty. Such utterances cannot directly be represented in Rolland's conceptual route graph. Our solution to this problem is the introduction of *place variables* and *orientation variables.* For example, *“irgendwo links oder rechts kommt dann dieser Turm”* (somewhere on the left or on the right appears this tower) is then represented as *ori(bx, p, right-neutral)* or *ori(bx, p, left-neutral)*, where “this tower” is at place *p*, the place variable *x* stands for the uncertain place on the path such that *p* is right or left to it. Such variables can often be resolved according to existing relations in the conceptual route graph and further route descriptions. In the case that a variable cannot be resolved**,** a subdialogue for requesting more information should be generated.

Granularity. Instructions occur on different levels of granularity, either referring directly to the goal, or listing route marks, or specifying each turn, etc. The level of granularity produced by any single user in the second part of the experiment does not necessarily match the level used in the previous part. If the conceptual route graph contains the information on the level of granularity employed by the user, then the instruction maps to the robot's knowledge. Problems arise in the two other kinds of possibilities. The utterance may contain reference to a destination that the robot does not know how to reach, or it may contain only low-level information that would lead to continuous need for interaction, in spite of the fact that the robot has access to higher level information in his knowledge representation. Thus, when a user only employs directional information such as “rechts” (right) or “geradeaus” (straightforward), corresponding route segments can not be created. Then, the communication situation would profit from a dialogue that induces users to employ the highest level of granularity that the robot is capable of dealing with. Furthermore, quantitative information such as “ungefähr zwanzig bis dreißig Sekunden” (about twenty to thirty seconds) – occurring in our data in the advance route description task – will be ignored if it is redundant, or will be replaced after the robot has received qualitative information via a subdialogue, since the conceptual route graph does not reflect quantitative knowledge.

Boundaries of tasks and actions. As stated in section 3 route segments directed from source place to target place are basic components required to build routes in the conceptual route graph. Thus, identifying boundaries of tasks and actions is important for completing route segments. However, many utterances do not contain explicit specification of spatial and temporal boundaries for tasks and actions, i.e., when or where do they start, and when or where are they completed. In some cases such boundaries are indicated in some way, either spatially or temporally, although – as exemplified above – such information may not always be useful to the robot. In other cases the boundaries need to be inferred, for example by investigating neighboring instructions. A spatial end boundary is indicated whenever people specify a goal or subgoal, as in “weiter geradeaus bis zum letzten Zimmer auf der linken Seite” (further straight on until the last room on the left side), or the length of a path is specified: “dreißig Meter vielleicht” (maybe 30 meters). Information about time is usually vague, as in “ziemlich lange geradeaus” (straight on for a fairly long time). In other cases an action ends when the next action starts, as in the spatial example “immer geradeaus, das nächste Mal wenn's links geht, links” (keep straight on, turn left the next time when it's possible to turn left); and the temporal one: “immer geradeaus, irgendwann links” (keep straight on, eventually left).

Temporal order. The above examples are indications that speakers make use of a temporal as well as spatial conceptualization, which can be more or less precise, and which may also differ widely between individual speakers. If the temporal sequence that the speakers have in mind for the completion of the task corresponds to the order of representation, then it can be directly represented as a route in the conceptual route graph, because a route is defined as a composition of a sequence of route segments with the temporal order. However, sometimes people go back mentally to add further information: "dann auf dem Flur nach links fahren bis zur ersten möglichen Abzweigung nach links, also da kommt dann noch dieser Turm aber der interessiert uns nicht wir wollen also nach links" (then drive on the corridor to the left until the first possible turn off to the left, I mean, there is also this tower but this does not interest us we want to go left). In this specific case the temporal digression is also a digression with respect to the task, i.e., the utterance contains further information about the scenario that is not directly related to the instruction (as indicated by "this does not interest us"). It is an open question how utterances should be dealt with that contain *instructions* in a distorted temporal order. We have not identified any such cases in our data, which may mean that it can be taken as the default case that instructions are given in the intended temporal order (suitable for reaching the goal). Digressions from this default will probably be indicated explicitly by linguistic means (such as "also" and "dann noch" in the above case). Using the conceptual route graph we can treat digressions as additional spatial relations on the route representation given in temporal order, but route instructions in a distorted temporal order can not be represented properly.

5. Dialogue modelling

According to the above analysis, in order to get a definitive route a dialogic interaction between the user and the robot is indispensable. Our approach to modelling dialogues is to move from empirically motivated dialogue patterns to models that can be used in computationally instantiated dialogue systems. We focus on the following concrete consequences that may come about due to one of the issues just listed, or a combination of several ones.

(1) The robot does not know some of the concepts contained in the user's instruction.

(2) The robot succeeds in identifying a direction of movement but not a destination. Thus, it does not know about the end boundary of the intended action.

(3) The robot does not know how to interpret the user's instruction because of some kind of underspecification.

(4) The robot recognizes that a destination has been mentioned but does not have access to information about how to get there, e.g. in the case of a granularity mismatch.

(5) The robot detects an instruction in a distorted temporal order.

(6) The robot finds that direction and destination, according to the instruction, do not match to each other.

In the first two cases, the robot should request the user to give more information. To treat problem (3) and (4), we can let the robot make some suggestion according to its knowledge, such that the user can give more precise information or use a more suitable level of granularity. If a conflict between the user's description and the robot's knowledge is detected (problem 6) or a temporally distorted instruction is detected (5), the robot should inform the user about it so that the user can correct the description if there is indeed a mistake, or choose a new way of giving the instruction.

The method followed here is based on the COnversational Roles (COR) model (Sitter & Stein 1992), which is a generic situation-independent dialogue-model. This model can be restricted or extended to cover precisely particular empirically motivated discourse patterns that are revealed during our experiments (Shi & Bateman 2005). Similar to *Information State* based approaches (e.g., Larsson & Traum 2000) the model accounts for particular abstract states of information that may occur in ongoing dialogue. Fig. 2 shows a recursive transition network representing a simplified dialogue model.

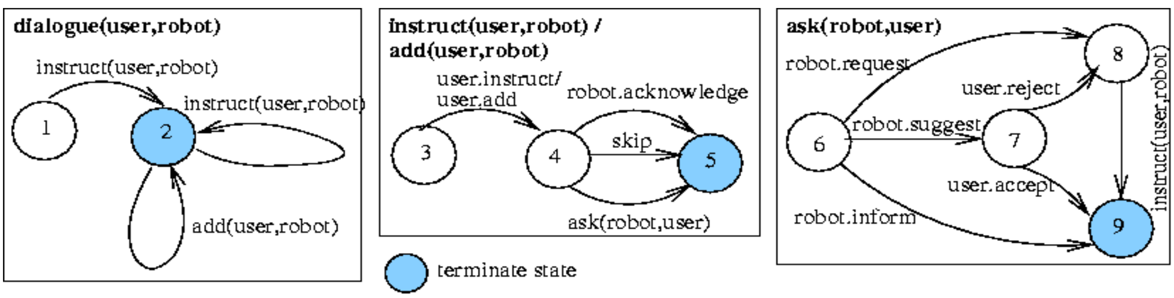

Fig. 2: The dialogue model as a transition network

Here the interaction possibilities have been split into three networks, the first describing the dialogic moves parameterized by participants, and the second giving the possibilities for the user to give instructions and additions, or for the robot to ask for extensive feedback. The third one focuses on the subdialogues generated by the robot to clarify potential problems that can arise due to the issues addressed in section 4. For modelling dialogues using transition networks we define the intentions of dialogue participants as a set of dialogue acts. Here we use *instruct* for giving instructions; *add* for giving additional information; *acknowledge* and *reject* for agreement and disagreement; *request*, *suggest* and *inform* for the robot to initiate subdialogues to clarify knowledge disparities and mismatches. In the transition network, parameterized moves, such as *instruct(user, robot)*, *add(user, robot)*, *dialogue(user, robot)* and *ask(robot, user)* are further defined as subdialogues. Moves like *user.instruct*, *user.add*, *robot.request*, and *robot.acknowledge* are basic dialogue moves; i.e., moves that are directly implemented by other components as speech acts and actions.

This dialogue model has been implemented in a prototypical dialogue system including a conceptual Route Graph covering the floor where the empirical data were collected. We are now working on the human language technology components to enable spoken language dialogue between users and Rolland in preparation of the next empirical round.

6. Conclusions

Our approach towards natural and efficient human-robot dialogues for route instruction tasks involves empirical investigation as well as dialogue modelling. Based on empirical results, we have presented a range of problematic aspects that can cause communication failure because of knowledge disparities between the robot's internal representation and the users' input, and suggested ways of addressing these problems via suitable and increasingly sophisticated clarification dialogues. Our experience shows that new problem classifications are needed if different knowledge representations or levels of knowledge abstractions are used, since diverging representations and abstractions are usually supported by different spatial reasoning methods. Moreover, with increasingly complex dialogue situations also the dialogue model presented here will need to be further expanded, since more expressive spatial calculi are necessary for representing and reasoning about users' knowledge, and further representation disparities may emerge. Further empirical research and analysis, increasingly involving real dialogic interaction, will gradually enable natural and efficient shared control between humans and robots in spatial tasks.

References

Bugmann, G., Klein E., Lauria S. & Kyriacou T. 2004. Corpus-Based Robotics: A Route Instruction Example. In *Proceedings of IAS*-8, 10-13 March 2004, Amsterdam, pp. 96-103.

Cohn, A. G. and Bennett, B. and Gooday, J. and Gotts, N. M. 1997. Qualitative Spatial Representation and Reasoning with the Region Connection Calculus. *Geoinformatics (1)*, Kluwer Academic Publishers, pp 1-44.

Denis, M. 1997. The Description of Routes: A Cognitive Approach to the Production of Spatial Discourse. *Cahiers de Psychologie Cognitive* (16), pp. 409-458.

Filipi, A. & R. Wales. 2004. Perspective-taking and Perspective-shifting as Socially Situated and Collaborative Actions. *Journal of Pragmatics, Volume 36, Issue 10.* October 2004, pp. 1851-1884.

Fischer, K. 2003. Linguistic Methods for Investigating Concepts in Use. In: Stolz, Th. and Kolbe, K. (eds.): *Methodologie in der Linguistik*. Frankfurt a.M.: Lang; pp 39-62.

Fischer, K. & R. Moratz. 2001. From Communicative Strategies to Cognitive Modelling. *Workshop Epigenetic Robotics*, Lund.

Freksa, C. 1992. Using Orientation Information for Qualitative Spatial Reasoning. In A. U. Frank, I. Campari, and U. Formentini (eds.), *Theories and Methods of Spatio-Temporal Reasoning in Geographic Space*. Berlin, Heidelberg: Springer, pp 162-178.

Krieg-Brückner, B., U. Frese, K. Lüttich, C. Mandel, T. Mossakowski, R. Ross. 2005. Specification of an Ontology for Route Graphs. In C. Freksa, M. Knauff, B. Krieg-Brückner, B. Nebel, and T. Barkowsky (eds.), *Spatial Cognition IV: Reasoning, Action, Interaction*. Berlin, Heidelberg: Springer, pp. 390-412.

Lankenau, A. & T. Röfer. 2000. The Role of Shared Control in Service Robots - The Bremen Autonomous Wheelchair as an Example. In: Röfer, T., Lankenau, A., Moratz, R. (Eds.): *Service Robotics -*

Applications and Safety Issues in an Emerging Market. Workshop Notes. European Conference on Artificial Intelligence 2000 (ECAI 2000), pp. 27-31.

Larsson, S. & D. Traum. 2000. Information State and Dialogue Management in the TRINDI Dialogue Move Engine Toolkit. In *Natural Language Engineering, Special Issue on Best Practice in Spoken Language Dialogue Systems Engineering*, pp. 323-340.

Ligozat, G. 2000. From Language to Motion, and Back: Generating and Using Route Descriptions. In: Christodoulakis, D.N. (ed.): *NLP 2000, LNCS 1835,* Berlin, Heidelberg: Springer, pp. 328-345.

Michon, P.E. & Denis, M. 2001. When and Why are Visual Landmarks Used in Giving Directions? In D.R. Montello (ed.): *Spatial Information Theory: Foundations of Geographic Information Science*. COSIT 2001 Morro Bay, CA, USA, September 19-23, 2001. Proceedings. Berlin, Heidelberg: Springer, pp. 292-305.

Moratz, R. & T. Tenbrink (in press). Spatial Reference in Linguistic Human-Robot Interaction: Iterative, Empirically Supported Development of a Model of Projective Relations. To be published in *Spatial Cognition and Computation*.

Shi, H. & Bateman, J. 2005. Developing Human-Robot Dialogue Management Formally. In *Proc. of the Symposium on Dialogue Modelling and Generation.* ISBN 3-921823-30-7, http://lubitsch.lili.uni-bielefeld.de/DMG/Proceedings/proc.html.

Sitter, S. & Stein, A. 1992. Modelling the Illocutionary Aspects of Information-Seeking Dialogues. *Information Processing and Management 28.* pp. 124-135.

Talmy, L. 2000. *Toward a Cognitive Semantics*, vol. I. Cambridge, MA: MIT Press.

Tenbrink, T. & Moratz, R. 2003. Group-based Spatial Reference in Linguistic Human-Robot Interaction. *Proc. EuroCogSci 2003: The European Cognitive Science Conference*, Osnabrück, Germany, pp. 325-330.

Tschander, L., Schmidtke, H., Habel, C., Eschenbach, C., Kulik, L. 2003. A Geometric Agent Following Route Instructions. In: Freksa, C., Brauer, W., Habel, C. & Wender, K.F. (Eds.), *Spatial Cognition III*: *Routes and Navigation, Human Memory and Learning, Spatial Representation and Spatial Learning*. Springer, pp. 89-111.

Tversky, B. 1999. Spatial Perspective in Descriptions. In P. Bloom, M.A. Peterson, L. Nadel & M.F. Garrett (eds.), *Language and Space*. Cambridge, MA: MIT Press, pp. 109-169.

Tversky, B. 2003. Structures of Mental Spaces – How People Think about Space. *Environment and Behavior* 35(1). Sage Publications, pp. 66-80.

Tversky, B. & Lee, P.U. 1998. How Space Structures Language. In Freksa, C., Habel, C. & Wender, K.F. (Eds), *Spatial Cognition: An Interdisciplinary Approach to Representation and Processing of Spatial Knowledge*. LNAI 1404, Springer, pp. 157-175.

Werner, S., B. Krieg-Brückner, & T. Herrmann. 2000. Modelling Navigational Knowledge by Route Graphs. In: Freksa, C., Brauer, W., Habel, C., & Wender K.F. (Eds). *Spatial Cognition II: Integrating Abstract Theories, Empirical Studies, Formal Methods, and Practical Applications.* Springer, pp. 295-316.