Spatial Strategies in Human-Robot Communication

Thora Tenbrink[1], Kerstin Fischer[2], Reinhard Moratz[3]

[1]University of Hamburg, Department for Informatics
Vogt-Kölln-Str. 30, 22527 Hamburg
tenbrink@informatik.uni-hamburg.de
[2]University of Bremen, Faculty for Linguistics and Literary Studies
Bibliothekstr. 1, 28359 Bremen, Germany
kerstinf@uni-bremen.de
[3]University of Bremen, Center for Computer Studies
Bibliothekstr. 1, 28359 Bremen, Germany
moratz@tzi.de

Abstract

This paper deals with various kinds of mental representations available for linguistic instruction in spatial human-robot interaction. After a survey of the literature on spatial representations in humans, focussing on dialogical interaction, we present a linguistic-pragmatic method for the analysis of spatial human-robot interaction. This method allows us to gain insights concerning the discourse strategies speakers employ, reflecting the underlying spatial representations in the reference systems employed. Furthermore, the users' hypotheses about the robot's functionalities are reflected in the succession of commands they try out in cases of failure. We describe an experiment involving human users and a robot that was designed on the basis of previous research on human spatial representations, and present new insights regarding both conceptual and spatial representational variation.

1 Mental Representations of Space and Discourse Strategies

How do human users communicate linguistically with robots about spatial configurations? The main difficulty in answering this question lies in the fact that most humans are, up to now, not used to talking to robots at all. While there is some evidence on how people deal with automatic dialogue systems (such as those employed for information services), we do not know much about people's strategies in talking to a robot in tasks involving a spatial environment. One reason for this is that most modern robots developed so far are not designed for linguistic communication. However, many kinds of contexts of human-robot interaction involve spatial configurations, such as the cooperative solving of spatial tasks. One very simple task of this kind is a scenario in which a robot is instructed to move towards a certain object. To indicate which object is meant, the instructor specifies its location. If such tasks are to be solved via linguistic communication, the question becomes crucial in which ways mental representations of space are used as resources for discourse strategies in the interaction between humans, robots, and the situative context.

Previous research in various areas provides us with useful insights as to what kinds of mental representations we should be prepared to deal with. In the wider field of *cognitive psychology*, one prominent area of research deals with *spatial reasoning*. In this area, one central question is what kinds of strategies humans employ when confronted with a difficult spatial problem, and what kinds of mental representations build the basis for these strategies. One way of solving such problems is to develop a so-called *mental model* [Johnson-Laird and Byrne, 1991], which serves to mentally visualize the problem at hand. A completely different strategy is to employ verbal propositions rather than imagery. Although a broad range of differing kinds of strategies has been identified already, the question which variables come into play in which kinds of situations is far from answered. One reason for this may be, as [Roberts and Newton, 2001] propose, that researchers often prefer simple tasks of 'high strategic purity' which, unfortunately, only offer themselves for a small range of possible strategies for problem solving.

A different approach to identifying human spatial representations is the analysis of linguistic phenomena. Psycholinguistic experimental studies on spatial situations focus on different kinds of mental representations that are reflected verbally through various kinds of *reference systems* [Herrmann and Grabowski, 1994]. In this area, valuable insights include the following:

- In localising reference objects in space, humans have - broadly speaking - three kinds of reference systems at their disposal, which (in Levinson's terminology [Levinson, 1996]) may be called *intrinsic*, *relative*, and *absolute*. In intrinsic reference systems, objects are located by referring to the intrinsic properties of another entity, such as the speaker's front in *The ball is in front of me*. Relative reference systems depend on the presence of a further entity (the so-called *relatum*), as in *The ball is in front of the table*. Absolute reference systems depend on the earth's cardinal directions, such as *north* or *south*.

- Additionally, speakers may variously employ either their own or their listener's *point of view* - or, which in some situations may also be useful, the perspective of a third entity (as in, *Viewed from the church's entrance, there is a bookshop on the right*). According to [Herrmann and Grabowski, 1994], speakers employ the listener's point of view specifically if there are reasons for this (such as a motivation to simplify the task for the listener, if the listener is a child or if the speaker wishes to be polite).

- In tasks involving route descriptions rather than the localisation of objects, some further kinds of perspectives are available to the speakers: for instance, one can assume the perspective of an 'imaginary wanderer' (an imagined person that walks along the route described). At the same time speakers may refer to landmarks available in the scenery [Herrmann and Grabowski, 1994].

- Speakers are not necessarily consistent in a discourse concerning the perspective they employ. According to [Tversky et al., 1999], speakers only choose objects as a landmark or as *relatum* if these objects are salient. Moreover, linguistically relatively simple reference terms seem to be preferred. This implies that speakers may choose to switch their perspective if this appears to be convenient in the current situation.

- A further problem is how objects are referred to linguistically in a discourse situation that allows a wide range of different conceptual representations. [Kessler et al., 1999], for example, show how reference resolution is achieved in dependence on the visual as well as the linguistic context; influenced, for instance, by the current focus of attention. They point to the importance of 'mental models' for reference.

- Finally, there is a vast body of research on the question which kinds of reference expressions are used by speakers with regard to different spatial areas. [Zimmer et al., 1998], for instance, show how simple expressions like *front* or *right* are used for straightforward spatial relations, while other relations call for more complex expressions such as *between front and left*. Other researchers test for the mechanisms of processing on linguistic, e.g. [Hörnig, 2001], and cognitive, e.g. [Wolff v., 2000], levels.

Obviously, the linguistic choices speakers make in interactions depend on the mental representations underlying their choices. But there is more to spatial interaction than the simple reflection of underlying spatial representations. In dialogue, speakers react to their interaction partner's contributions, and they attune their linguistic choices to what they believe to be suitable for their partner in the situation at hand. For example, [Schegloff, 1972] shows how 'formulating place' depends on the recipient for whom the description is designed. [Schober, 1993] found that speakers attend to their hearers' clues as to whether they have understood the instruction in the sense that the references have been *grounded* [Clark and Wilkes-Gibbs, 1986]. [Garrod and Anderson, 1987] found that communication partners interactively developed distinct but consistent description schemes, which reflected different kinds of underlying mental representations (which, again, may be labeled mental models). Such representations were very much dependent on the interaction itself as well as on the given task (see also [Rieser, 1996] for related findings).

Several research groups in Europe are currently concerned with different aspects of spatial conversation in human-robot interaction. In [Ligozat, 2000] achievements of the LIMSI research group in Orsay (Paris) are described, dealing with spatial robot instruction as well as output generation in spatial tasks. The Collaborative Research Center (SFB) 378 "Resource-adaptive cognitive processes" in Saarbrücken deals, among other things, with the question how human-robot dialogues can be designed to adapt to the users' cognitive resources. Related to this work (as part of the earlier SFB 314 in Saarbrücken), [Stopp et al., 1994] considered many of the above-mentioned factors involved in spatial reference in their approach to spatial human-robot interaction. They point out that speakers do not specify every detail needed for unambigous reference in instructing robots. Furthermore, the SFB 360 "Situated Artificial Communicators" in Bielefeld deals with the question how situated and integrated communication can be achieved effectively in a robot instruction task involving the construction of a toy air plane [Moratz et al., 1995], allowing dialogic communication between human and robot.

However, the specific effect of a *robot* as an interaction partner on the linguistic and spatial choices of a human speaker has not been addressed so far. As previous studies in the related field of human-computer interaction, e.g. [Amalberti et al., 1993, Fischer, 2000], have shown, the users' conceptualisation of their interaction partners has considerable impact on their speech. Such effects also need to be worked out for human-robot interaction.

Concluding from the discussion so far, a number of questions are still open:

- How does the choice of discourse strategies used in interaction reflect the mental representations available to the participants?
- What factors determine the speaker's choice of particular discourse strategies specifically in human-robot interaction, such as specific reference systems, perspectives, reference objects or landmarks?
- What kinds of dialogic aspects influence the speakers' strategies?
- How is the speaker's conceptualisation of the robot as an interaction partner reflected in the discourse strategies taken?

To answer these questions, it is necessary to experiment with users unfamiliar with the technological and linguistic properties of their interaction partner, the robot, rather than focusing on the operativeness of a newly developed system.

In the remainder of this paper, we outline some central aspects of the method we employ in our approach. To exemplify the method proposed, we present the results of our first exploratory study, addressing some of these issues.

2 A Linguistic-Pragmatic Method for Analysing Human-Robot Communication

The methodology presented here builds on perspectives developed in the framework of conversation analysis (CA). Two central notions of CA are especially relevant to our aims: First, a central issue is the concentration on those aspects of the communicative situation of which it can be shown that they are relevant for speakers when deciding on a particular strategy. CA provides methods for verifying those relevant aspects, instead of claiming *a priori* the importance of specific situational factors. Second, in CA the notion of *deviant case analysis* [Hutchby and Wooffitt, 1998] is central: cases of miscommunication are especially suitable for the analysis of underlying speaker strategies. Human-robot interaction can be viewed as deviant in comparison to natural human-to-human communication. Miscommunication is especially likely to occur between a newly developed robot system and a user that is unfamiliar with it. Moreover, crucial aspects of the communicative situation, such as the spatial setting, the definition of the situation as formal or informal, the robot's linguistic properties, or the appearance of the robot, can be manipulated in the study without necessarily increasing the unnaturalness of the already unnatural situation. This allows the controlled investigation of a number of variables that crucially influence communicative processes both in human-robot interaction and human-to-human communication (such as recipient design, alignment, interactive negotiation, or the role of linguistic feedback). These would be much harder to control in natural human-to-human communication.

Human-robot interaction also provides us with a number of additional data not usually available in human-to-human communication. Users often produce self-talk in which they give accounts of their strategies, and in which they reveal their interpretations and explanations about what is going on. For example, speakers may overtly announce discourse strategies such as the repetition of previously produced utterances in cases of miscommunication [Fischer, 1999, Fischer, 2000]. Besides the naturally occuring self-talk, it is also possible to ask the participants to 'think aloud', i.e., to verbalise their strategies and their reflections on the discourse situation, or to fill out a questionnaire after the conversation with the system. These data can provide insights with regard to which kinds of information speakers attend, and which factors influence their choice of strategies.

The procedure proposed here is thus the following: A body of linguistic instances of spatial human-robot communication is collected experimentally. Linguistic analysis of the data then reveals different kinds of reference systems used in specific kinds of tasks, different kinds of underlying representations, and various interactional strategies which may depend both on the specific situation and on the robot's reactions. Furthermore, the temporal order in which discourse strategies are employed by the speakers may reveal their hypotheses about spatial instruction. A pragmatic analysis thus both provides an overview of the strategies speakers employ in spatial instruction and points to the factors that influence strategy selection.

3 Example Analysis: Strategy Selection in Human-Robot Communication

In our first study involving experiments with a robot prototype, our aim was to analyse the way human users interact with the robot, which spatial strategies they employ, and to which aspects of the situation they attend for selecting their strategies. A further aim was to determine how these strategies are adapted during the interaction with the system.

3.1 The Robot

The robot system uses a Pioneer I as mobile basis (see figure 1). The robot is equipped with an elevated camera, which observes the scene in front of the robot with a wide angle lense. A colour segmentation module delivers objects of high colour saturation which are categorised into given classes of objects. Using the internal model of the observed scene, the robot can plan paths that avoid obstacles. Natural language input via a keyboard is analysed using a Categorial Grammar parser [Hildebrandt and Eikmeyer, 1999]. The referenced goal is matched with a position in the perceived scene, a path is planned and the movement is executed. Details of the experimental system and the architecture are described in [Moratz and Fischer, 2000, Habel et al., 1999].

3.2 Experimental Design

A test scenario was developed in which the user's task was to make the robot move towards particular objects pointed at by the experimentator. Users were asked to type natural language sentences into a computer in order to instruct the robot. The setting was a joint attention scenario in which a number of cubes were placed on the floor together with the robot, for instance, in a 90 degree angle or opposite of the participant, as shown in figure 2. The actual arrangements of the cubes was varied; in one quarter of the settings, a cardboard box was furthermore added to the setting in order to trigger instructions referring to the box as a salient object.

The robot was designed to process qualitative linguistic information such as "go to the block on the right". If

Figure 1: Our experimental platform, the mobile robot GIRAFFE

a command was successful, the robot moved to the block it identified. The only other possible response was “error”. Thus, users who were not successful from the start were challenged to try out many different kinds of spatial instruction to enable the robot to identify the intended aim, without being prompted to a particular solution by the robot.

15 different participants carried out an average of 30 attempts to move the robot within about 30 minutes time each. Their sentences were protocolled, and their verbal behaviour during the experiments was recorded in order to capture self-talk. After the experiments, participants filled in questionnaires dealing with their strategies in communicating with the robot.

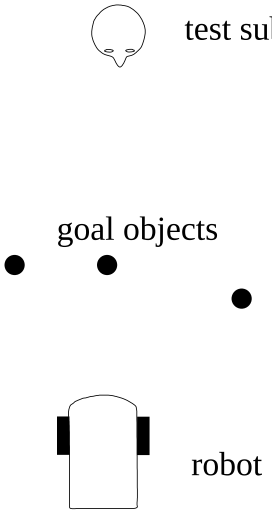

Figure 2: The setting of the experiment

3.3 Experimental Results: Instructional Strategy Selection

Our linguistic-pragmatic analysis revealed new insights particularly in two areas: on the one hand, we identified a range of variations in the users’ *choice of spatial reference systems*; on the other, some crucial aspects of the *conceptual representation* underlying the instructional strategies were revealed. Note that this variability, which we will work out in more detail below[1], is specific to the given interaction situation, which (like many typical scenarios in human-robot interaction) involved a human user, a robot, and the task of moving to one particular object.

Concerning the first type of results we found that speakers made extensive use of the concept of a *group of similar objects* for spatial reference. This concept has been largely ignored in the literature on object localisation, which may in part be due to differences in the experimental setting. In our setting, the participants were not - as in typical psycholinguistic experiments - asked to specify the location of one object in relation to a different one, but rather, to specify the identity of one of several similar objects whose location was known. Unambiguous reference could then be achieved effectively by naming

[1]Please also refer to [Moratz et al., 2002] for further details

the object's position relative to the rest of the group of similar objects. Many users employed this strategy, using linguistic instructions such as *fahr zum rechten Würfel* [drive to the rightmost cube]. Thus, the speakers' instructions revealed their mental representation of the objects as a group. In other types of instruction, users oriented at the robot's intrinsic properties instead. In this case, the users' mental representation of the salience of body orientation for view or movement directions was revealed, as in *Fahr zum Würfel rechts von dir* [Drive to the cube to your right].

The second type of results shows a consistent order in which different kinds of instructions were employed. In particular, only half of the participants started their attempts to instruct the robot by naming the intended reference object itself, as in the above examples. This strategy was the one we expected and implemented, with the effect that these instructions were usually successful, unless there were orthographic, lexical, or syntactic problems. In such cases, these participants used directional strategies in later instructions; if successful, they reapplied the goal-naming strategy. The other half of the participants started by giving directional descriptions. This implied a decomposition of the main action into smaller portions, as illustrated by instructions such as *fahr 1 Meter geradeaus* [drive 1 meter ahead], or *rolle ein wenig nach vorn* [roll a bit forward]. If the direction descriptions did not work, the participants did not try out a description of the goal object, which the robot would have understood. Instead, they used descriptions of movements that were unspecified regarding a particular direction, for instance *fahre* [drive], *Drehung!* [turn!]. Some participants who had used this strategy employed afterwards a fourth one, namely to specify the instrumental actions necessary for such movement, for example: *drehe Deine hinteren Rollen* [turn your rear wheels] or *Motor an* [engine on].

Thus, the order of instructions employed by the users revealed the following hierarchy of instructional strategies:

	goal description
$<$	direction description
$<$	movement description
$<$	description of actions instrumental to movement

We propose that this consistent order of instructional strategies reflects the participants' hypotheses of the domain of spatial instruction: namely, that they regard knowledge about how to move into a specific direction instrumental to moving towards a goal object, that they regard knowing how to move at all instrumental to moving into a specific direction, and that they consider knowing about how to use one's facilities for moving instrumental for moving. Moreover, participants seemed to ignore the possibility that a robot could know how to move to a goal object without being able to *understand* directional instructions.

4 Factors that May Influence Strategy Selection

According to our analysis, speakers order their strategies in the way they do because of their hypotheses about *basicness and difficulty*. In particular, those speakers who did not try out the goal naming strategy at all may have assumed that this kind of complex instruction is too difficult for the robot. In the following, we look for further evidence that supports our hypothesis that in this particular situation, basicness and difficulty is relevant for the speakers. There are several observations that point in the same direction:

1. **Point of View**

   Unlike in communication among humans [Herrmann and Grabowski, 1994], the speakers in our experiment *consistently* took the robot's perspective, unless there was (or seemed to be) evidence that this could not be the right strategy. This linguistic behaviour may indicate that the speakers regarded the robot as a communication partner who is not capable of taking the speaker's perspective, i.e., who should receive as simple instructions as possible.

2. **Group-based Reference**

   As pointed out above, many participants made use of the concept of a *group* in order to specify the position of one of its members. However, the question needs to be asked why many users did *not* use this concept, as it turned out to offer (in this scenario) an unambiguous referential strategy involving a linguistically simple kind of instruction. One reason for many users' failure to take advantage of this might be that the users did not expect the robot to be able to grasp the concept of a group, as this involves comparison, identification of similarity, and categorisation.

3. **Linguistic Constructions**

   Speakers wondered both during the experiments and in the questionnaires about the linguistic capabilities of the robot, asking whether it understood particular words or syntactic constructions, such as relative clauses. Thus, they attended to the fact that the robot might have limited linguistic capabilities. Furthermore, most speakers employed jussive imperatives, a linguistic strategy rarely used in task-oriented human-to-human dialogues, as it completely lacks the various (sometimes rather complex) kinds of elaborations which are considered polite.

Taken together, these findings suggest that the interaction in our experiments was influenced by the speakers' conceptualisation of the robot as a communication partner with non-humanlike capabilities.

5 Conclusions and Prospects

In this paper, we have presented a method for investigating the selection of particular instructional strategies in human-robot interaction. The results of our study show strategy variation with regard to conceptual representations and spatial reference systems. We showed that the particular choice of an instructional strategy is influenced, for instance, by the speakers' conceptualisation of the robot as a communication partner who needs comparably simple instructions. Furthermore, our results point to the users' consideration of various kinds of spatial representation. Thus, this kind of analysis contributes to bridging the gap between results achieved in various research areas on mental representations of space, and the actual usage of strategies for spatial instruction observable in human-robot interaction.

So far, we have addressed only a small portion of the issues that we expect to be relevant for spatial human-robot interaction. Recall that previous experiments in other research areas revealed that much variation is to be expected dependent on the spatial setting, the specific scenario employed in the experiment, and various other factors as listed above. Therefore, it is necessary to consider many different kinds of settings in which it becomes possible to explore a greater range of varieties of user conceptualisations of spatial configurations, of the robot's functionalities, and so on. Moreover, there are good reasons to assume that verbal responses by the robot would have great influence on the users' strategies [Zoltan-Ford, 1991]. In future research, we will therefore also explore the ways in which the dialogue itself contributes to the users' decisions on which mental representations will be most suitable for the spatial task at hand.

References

[Amalberti et al., 1993] Amalberti, R., Carbonell, N., and Falzon, P. (1993). User Representations of Computer Systems in Human–Computer speech interaction . *International Journal of Man–Machine Studies*, 38:547–566.

[Clark and Wilkes-Gibbs, 1986] Clark, H. and Wilkes-Gibbs, D. (1986). Referring as a collaborative process. *Cognition*, 22:1–39.

[Fischer, 1999] Fischer, K. (1999). Repeats, reformulations, and emotional speech: Evidence for the design of human-computer speech interfaces. In Bullinger, H.-J. and Ziegler, J., editors, *Human-Computer Interaction: Ergonomics and User Interfaces, Volume 1 of the Proceedings of the 8th International Conference on Human-Computer Interaction, Munich, Germany.*, pages 560–565. Lawrence Erlbaum Ass., London.

[Fischer, 2000] Fischer, K. (2000). What is a situation? *Proceedings of Götalog 2000, Fourth Workshop on the Semantics and Pragmatics of Dialogue, Göteborg University, 15-17 June 2000. Gothenburg Papers in Computational Linguistics*, 00(05):85–92.

[Garrod and Anderson, 1987] Garrod, S. and Anderson, A. (1987). Saying what you mean in dialogue: A study in conceptual and semantic co-ordination. *Cognition*, 27:181–218.

[Habel et al., 1999] Habel, C., Hildebrandt, B., and Moratz, R. (1999). Interactive robot navigation based on qualitative spatial representations. In Wachsmuth, I. and Jung, B., editors, *Proceedings Kogwis99*, pages 219–225, St. Augustin. infix.

[Herrmann and Grabowski, 1994] Herrmann, T. and Grabowski, J. (1994). *Sprechen: Psychologie der Sprachproduktion*. Spektrum Verlag, Heidelberg.

[Hildebrandt and Eikmeyer, 1999] Hildebrandt, B. and Eikmeyer, H.-J. (1999). *"Sprachverarbeitung mit Combinatory Categorial Grammar: Inkrementalität & Effizienz "*. SFB 360: Situierte Künstliche Kommunikatoren, Report 99/05, Bielefeld.

[Hörnig, 2001] Hörnig, R. (2001). Are ambiguous anaphors assigned in such a way to avoid ambiguous spatial descriptions? In *Proceedings KogWis 2001 5th Meeting of the GermanCognitive Science Society*.

[Hutchby and Wooffitt, 1998] Hutchby, I. and Wooffitt, R. (1998). *Conversation Analysis*. Cambridge: Polity.

[Johnson-Laird and Byrne, 1991] Johnson-Laird, P. and Byrne, R. (1991). *Deduction*. Hove: Psychology Press.

[Kessler et al., 1999] Kessler, K., Duwe, I., and Strohner, H. (1999). Subconceptual Dynamics in the Resolution of Reference in Discourse. In Rickheit, G. and Habel, C., editors, *Mental Models in Discourse Processing and Reasoning.* Elsevier, Amsterdam.

[Levinson, 1996] Levinson, S. C. (1996). Frames of Reference and Molyneux"s Question: Crosslinguistic Evidence . In Bloom, P., Peterson, M., Nadel, L., and Garrett, M., editors, *Language and Space*, pages 109–169. MIT Press, Cambridge, MA.

[Ligozat, 2000] Ligozat, G. (2000). From language to motion, and back: Generating and using route descriptions. In *Natural Language Processing 2000*.

[Moratz et al., 1995] Moratz, R., Eikmeyer, H., Hildebrandt, B., Knoll, A., Kummert, F., Rickheit, G., and Sagerer, G. (1995). Selective visual perception driven by cues from speech processing. In Pinto-Ferreira, C. and Mamede, N., editors, *Advanced Manufacturing Forum, Vol. 1*. Trans Tech Publications, Schweiz.

[Moratz and Fischer, 2000] Moratz, R. and Fischer, K. (2000). Cognitively adequate modelling of spatial reference in human-robotinteraction. In *IEEE International Conference on Tools with Artificial Intelligence (ICTAI-2000)*, pages 222 – 228, Vancouver, British Columbia, Canada.

[Moratz et al., 2002] Moratz, R., Tenbrink, T., Fischer, F., and Bateman, J. (2002). Spatial knowledge representation for human-robot interaction. In Freksa, C., Brauer, W., Habel., C., and Wender, K., editors, *Spatial Cognition III*. Springer, Berlin. *in print* .

[Rieser, 1996] Rieser, H. (1996). Repräsentationsmetonymie, perspektive und koordination in aufgabenorientierten dialogen. In Umbach, C., Grabski, M., and Hörnig, R., editors, *Perspektive in Sprache und Raum*. Deutscher Universitäts-Verlag, Wiesbaden.

[Roberts and Newton, 2001] Roberts, M. J. and Newton, E. J. (2001). Understanding strategy selection. *International Journal of Human-Computer Studies*, 54:137–154.

[Schegloff, 1972] Schegloff, E. A. (1972). Notes on a conversational practice: Formulating place. In Sudnow, D., editor, *Studies in Social Interaction*, pages 75–119. New York: The Free Press.

[Schober, 1993] Schober, M. (1993). Spatial perspective-taking in conversation. *Cognition*, 47:1–24.

[Stopp et al., 1994] Stopp, E., Gapp, K., Herzog, G., Laengle, T., and Lueth, T. (1994). *Utilizing Spatial Relations for Natural Language Access to an Autonomous Mobile Robot*. Sonderforschungsbereich 314, Bericht 107, Saarbrücken.

[Tversky et al., 1999] Tversky, B., Lee, P., and Mainwaring, S. (1999). Why do speakers mix perspectives? *Spatial Cognition and Computation*, 1 (4):399–412.

[Wolff v., 2000] Wolff v., A. (2000). Ego-reorientation on mental representations of rooms. *Cognitive Processing, International Quarterly of Cognitive Science*, 1.

[Zimmer et al., 1998] Zimmer, H., Speiser, H. R., Blocher, A., and Stopp, E. (1998). The Use of Locative Expressions in Dependence of the Spatial Relation between Target and Reference Object in Two-Dimensional Layouts. In Freksa, C., Habel, C., and Wender, K., editors, *Spatial Cognition. An Interdisciplinary Approach to Representing and Processing Spatial Knowledge*, Lecture Notes in Artificial Intelligence, pages 223–240. Springer-Verlag, Berlin.

[Zoltan-Ford, 1991] Zoltan-Ford, E. (1991). How to get people to say and type what computers can understand. *International Journal of Man-Machine Studies*, 34:527–647.