

Methods for analyzing natural discourse: Investigating spatial language in HRI vs. in a no-feedback web study

Thora Tenbrink

Universität Bremen, SFB/TR 8 Spatial Cognition
tenbrink@sftr8.uni-bremen.de

Abstract. The focus of interest in my research lies in the investigation of spontaneously produced natural language used to refer to the spatial position of a goal object. In this short paper I compare two central elicitation scenarios which have been useful for the investigation of speakers' strategies to achieve given discourse purposes by using spatial reference: a no-feedback web study and a human-robot interaction scenario. In both cases the task was to identify one out of several similar objects in a configuration by using spatial reference. The results of the two kinds of studies show a number of important systematic differences as well as striking parallels with respect to speakers' conceptual and linguistic strategies.

Introduction

Spatial reference (i.e., utterances such as “the object on the left”) is a central issue within the wider field of Spatial Cognition because it encompasses a number of central aspects relevant to diverse disciplines. First, since spatial reference is used to represent a speaker's conceptualization of a spatial scene in natural language, the researcher gains access to underlying cognitive processes by analyzing the linguistic product. Second, spatial reference occurs as a subtask in a number of cognitively demanding tasks such as realizing, identifying, and communicating the position of objects, navigation, wayfinding and route communication, scene description and representation, architectural design, or object assembly and construction. Many of these tasks have been subject to psychological and related studies. Finally, the wider field of object reference is one of the central issues within linguistic pragmatics with a wide range of related empirical study, formalization, and realization in automatic systems.

Given the broad range of interest, it comes as a surprise that there are still a number of facets within the field of spatial reference that have barely been touched by empirical research. In fact, most investigation of spatial language can be accredited to psycholinguistics, leading to a wealth of information with respect to language comprehension (including acceptance ratings) in meticulously controlled settings, and limited production (typically restricted to pre-defined syntactic formats). This approach has its undisputed merits; it has led to the identification of many central as-

pects of the spatial setting, the previous discourse, underlying concepts, and processing issues that influence the production and comprehension of spatial language terms in systematic ways (for a detailed overview of psycholinguistic and related results concerning spatial projective terms see Tenbrink 2005d). Unfortunately, since free production studies are rare, it is not known whether, to what degree, or under which exact circumstances do the types of structures that are tested in comprehension studies coincide with those that are used spontaneously by speakers. Since only few attempts have so far been made to combine such research efforts with approaches and methods from linguistic pragmatics (cf. Sperber & Noveck 2006), such as discourse analysis, the relation to speakers' spontaneous productions in natural discourse (produced in natural settings) has not been established sufficiently. We therefore do not know much about how the findings from psycholinguistic laboratory work might fit with the insights gained from the analysis of natural discourse; in most cases, this is true to such a high degree that instances of spatial terms occurring in larger corpora (collected without restricting speakers' linguistic choices) cannot be interpreted sufficiently in the light of earlier results, as too many factors still remain unaddressed.

Of course, there are notable exceptions to this, encouraging us to work further on establishing the bridge between control and freedom of choice. Many far-reaching ideas and detailed results stem from recent investigations of natural discourse (e.g., Clark 1996, Schober & Brennan 2003, Pickering and Garrod 2004); quite often in such work, dialogues were collected that relate to a spatial setting (see also the now completed Maptask project in Edinburgh, Anderson et al. 1991). Also, the field of object reference is maturing towards generality by the identification of systematic discourse processes leading to the preference of one linguistic strategy over another in relation to the given setting (e.g., Herrmann & Deutsch 1976, Dale & Reiter 1995, van Deemter & Kibble 2002). Nevertheless, a broad range of aspects and discourse tasks remain for which the relationship between detailed findings and spontaneous discourse still needs to be established, for example, by increasing naturalness of experimental settings, and comparisons with qualitative corpora analyses. Methods for such approaches are currently being established (cf. Sperber & Noveck 2006; Fischer 2003, de Vega & Rodrigo 2005).

The point of the present work is to contribute to this issue by exploring the feasibility of controlling the (spatial and discourse-related) setting and discourse task but not the language to be used to achieve the given aim, in order to create a reliable source for the analysis of spontaneously used natural language. As an example for this, the specific discourse task of contrastive spatial reference is focused on. The question of how one object is naturally referred to in a setting containing several objects of the same class by using spatial reference has seldom been addressed in earlier work (though see van der Sluis & Krahmer 2000, Gorniak & Roy 2004). To investigate speakers' spontaneous linguistic choices, I analyzed the language data from two different corpora collected in experimental studies in line with the present discourse-analytic approach in which linguistic strategies are in focus. The two discourse settings employed here are a no-feedback web study and a human-robot interaction study. They both concern a simple and restricted task of contrastive spatial reference, and they basically relate to the same set of spatial configurations. Although the settings themselves are not natural per se, at least they provide a setting in which speakers can freely produce language and develop their own strategies towards solving the

given discourse task. In contrast to a corpora analysis that is purely based on naturally occurring texts, this approach ensures that the influence of the (spatial) context can be accounted for adequately, which is important because of the particularly high context dependency of spatial terms. This is therefore considered as a step towards the analysis of natural discourse. In the following, I will briefly present and compare the results with respect to the spontaneous application of spatial terms in this discourse task.

Experimental Studies

Study I: Web study




Fig. 1. One picture used in the web study.

The web study (Tenbrink 2005a,b) served to elicit participants' reference to a cued object on the basis of its spatial position. The participants (native speakers of English or German) were confronted with 14 different pictures (such as the one depicted in **Fig. 1**) in a sequence, containing two or three objects of the same class (squares), sometimes together with another object of a different class (a circle), a speaker position inside the picture (marked by an "X"), and an addressee position (marked by an "Y"). The participants were asked to single out the marked object, for example, by imagining themselves in position X and instructing person Y to go to the object marked with a circle. No syntactic format was prescribed, and no example was given. Therefore, the scenario and discourse task, though simplistic in nature, allowed for a broad variability with respect to choices on all linguistic levels: this includes diverse grammatical constructions, variability in information structure, and lexical choice, as well as conceptual options for reference systems, perspectives, and spatial strategies such as describing the path towards a goal vs. directly referring to the goal object's position. Since there was no feedback whatsoever, speakers relied solely on their own intuitions about how to reach the given discourse aim.

Altogether, a total of 2,332 utterances produced by native German speakers and 1,480 utterances produced by native English speakers were analyzed. The data were examined with respect to choice of strategy, alternatives to projective terms in goal based instructions, and details about the application of projective terms, including linguistic variability and conceptual choices such as perspective, relatum, and axes, which together serve as the basis for a reference system (see Tenbrink 2005c for details). However, the reference system itself cannot be identified directly on the basis of the linguistic form. This is due, on the one hand, to an overwhelming lack of ex-

plicitness with respect to either relatum or perspective, or both, and on the other hand, to the fact that there is no one-to-one correspondence between linguistic forms and underlying reference systems. If both elements are mentioned explicitly, the reference system can be identified. But as the data show, perspectives are almost never (specifically in English) given explicitly, and relata are mentioned only in a subset of cases, typically to avoid obvious misunderstandings. Therefore, the data point to the conclusion that in the majority of situations there is more than one underlying reference system that is compatible with the linguistic representation in relation to the situation at hand. Further specifications and restrictions are provided by the speaker only if they are necessary for discrimination.

Generally, participants use a broad spectrum of variability on all scales. The analysis shows that linguistic as well as conceptual choices depend heavily on the spatial situation, i.e., the presence of other objects and (imagined) persons, and the available kinds of perspective. A range of systematic strategies and speakers' preferences in relation to the situational factors were specified (Tenbrink 2005b). Additionally, a number of language-specific differences could be identified (see Tenbrink 2005a).

Study II: Human-Robot Interaction

The participants in our human-robot interaction study were confronted with a scenario that contained similar spatial configurations (object arrangements) as those depicted in the web study, this time in a real world situation (cf. Moratz & Tenbrink 2006). Large cardboard boxes were positioned at the locations of the squares in the pictures, a barrel replaced the circle, and the participant and the robot were positioned at the places of X and Y of the web study. The task was to instruct the robot to move to a cued goal object. In this situation, the robot could only carry out those instructions that it was equipped to understand, namely, those containing goal-based spatial reference (such as "go to the box on the left"; in German). If the utterance could be interpreted, the robot moved to the goal object. Otherwise, there was only brief verbal feedback. In the first part of the study, the robot could only say "I don't understand"; in the second part, it produced more differentiated feedback based on the module that failed to understand the given input: thus, the user was told if the robot's lexicon did not contain one of the words, or if the parser could not understand the syntactic structure, or if the computational model could not interpret the spatial description with respect to the perceived situation. Also, the robot sometimes produced scene descriptions in order to give the user a cue as to its functionalities. Therefore, speakers' utterances were based first on their own intuitions as in the web study, but increasingly, speakers could build on their own experience with respect to robotic output and feedback as well as success and failure.

Since this is the case, analysis along precisely the same lines as in the web study was not feasible, since the discourse processes involved differ in fundamental respects. Therefore, the utterances produced in the HRI study were analyzed in relation to the robot's capabilities. In the first part of the study, 605 written and 778 spoken instructions to the robot were collected, and in the second part of the study, 482 written instructions were collected.

It turned out that users frequently rely on syntactic structures that had previously proved successful. Thus, they take previous success into account in a similar way as human interactants react to positive feedback, which is usually in that case given verbally or by paralinguistic signaling. In subsequent discourse, the successful description format (usually a projective adjective) is further employed even if it does not conform to the usual applicability preferences in a different configuration, adapting only the specific projective term according to the given target object (e.g., exchanging *left* for *right* etc.). In general, uninformed users are extremely creative in their usage of spatial reference. Even if informed about the basic features of the robot, they introduce new and interesting ways of referring to the goal object, especially after communicative failure or in cases of difficult referability.

The differentiated robotic output produced in the second part of the study led to a major increase in the success rate. With respect to the written instructions of the first part of the study, users needed an average of four turns to achieve success in their instruction, and they only managed to successfully instruct the robot in 64% of the tasks. In the second part of the study, in contrast, 95% of tasks were successfully solved after two user instructions on average. Clearly, the users' intuitive hypotheses about the robot's abilities were wrong to such a high degree that they needed substantially more feedback than was provided by the robot's behavior. This is particularly due to the fact that users typically did not simply modify the linguistic aspects of their spatial description after failure, but they assumed that their description was not simple enough and therefore switched to a lower level of instruction (cf. Fischer & Moratz 2001; see below).

The spoken utterances (in the first part of the study) produced additional complications due to the present state of the art of speech recognition. Importantly, in cases of communication failure speakers immediately assumed that their instructions were too difficult for the robot; therefore, goal based instructions (which would have been successful in written language) were not repeated. Thus, a single case of speech recognition failure could have enduring consequences for the ensuing discourse (Moratz & Tenbrink 2003).

Spatial reference in both settings: General results

Both kinds of studies can be considered as unnatural in some respects: the web study is certainly unnatural because participants had to imagine themselves in a real world situation, partly including an interaction partner; the situation depicted was in 2D, representing a 3D setting; and the scenario lacked embedding in a suitable social situation. The HRI study, although situated in the real world, is a peculiar situation in itself, given that most humans do not have any experience at all in talking to robots. Furthermore, the discourse in the HRI studies depended heavily on the way the robot was designed, i.e., on the system designers' expectations about what users would produce: if speakers happened to use a "correct", i.e., conforming, kind of utterance, they were successful; otherwise, they had to try out alternatives. The fact that the two kinds of studies differ profoundly with respect to each other, including the problems associated with them, does not improve the situation.

In spite of these obvious drawbacks, the results of the two studies highlight a number of important systematic differences as well as striking parallels with respect to speakers' conceptual and linguistic strategies. For example, although the unnatural situations seemed to induce participants to use other conceptual and linguistic strategies than expected, those utterances that did contain spatial reference on the basis of projective terms – which was the main aim of the studies – were remarkably consistent. The results point to underlying principles of reference that are in accord with previous studies in several ways. In the following, I will work out systematic results and general conclusions that can be drawn in spite of the limitations of the present approach.

Altogether, in both cases, the speakers' utterances roughly fell into the following categories: only-path descriptions (basically, route descriptions towards the goal without mentioning the goal object itself, such as *go to the left*), path-plus-goal descriptions (route descriptions including the endpoint, such as *move forward to the box*), and goal-based descriptions (identifying the endpoint only by using spatial reference, such as *go to the box on the left*). The main difference between the two scenarios concerns the relative frequency of utterances within these three categories, also with regard to changes with time throughout the studies. Most inexperienced users in the HRI study did not mention the goal object at first and then had to gain a considerable amount of experience with the robot in order to change this basic strategy, as the robot could not understand these utterances. In contrast, most participants in the web study produced goal-based descriptions from the start, contrasting the goal object from the competing ones by suitable spatial descriptions such as *the leftmost square*.

Some of these speaker strategies, such as employing detailed route descriptions to describe the goal object's position, were not expected prior to the studies. Typically, route descriptions relate to large-scale environments rather than indoor scenarios or those that can be perceived at a glance, which are sometimes referred to as **vista space** (Montello 1993). However, there are good reasons for this result with respect to both kinds of studies: the participants in the web study were not informed about scale, i.e., the size of the objects and the distances between them were unknown. Thus, although they could indeed perceive the whole scene in the study itself, they may have conceptualized the scene as a large-scale environment. This is plausible because the study required a lot of imagination from the start. The participants in the HRI study, on the other hand, were not informed about the robot's perceptual abilities. Thus, they may have assumed that the robot was not capable of comprehending the scene as such; therefore, they described the movements necessary to reach the goal. In addition to this, in the HRI tasks speakers sometimes resorted to lower level descriptions, describing preliminary actions that do not directly relate to the spatial task at hand (see Fischer & Moratz 2001).

In the analysis of the goal-based utterances, further systematic results could be found concerning the choice of perspectives and reference systems in this kind of discourse task. With respect to perspective, results are straightforward and seem to be largely independent of the spatial configuration: If a human (or human-like entity, such as a robot) is present, their position usually serves as origin, the basis for perspective. In a dialogic situation, typically the interlocutor's perspective is taken. Here, the principle of **partner adaptation** is crucial. Perhaps because of the general agreement about the perspective to be adopted, perspectives are seldom given explicitly.

With respect to reference systems, results are more complex. As mentioned above, in many cases the description was consistent with more than one underlying reference system, as some of the ingredients were not mentioned explicitly. Among the reference systems that could be identified, intrinsic systems are the most frequent, using the preferred origin also as relatum. However, in German there is a second common option, which does not seem to be as frequently employed in English, namely, using a group of similar objects as a relatum. Other relata were much less frequently used throughout the studies. The choice of relatum may be influenced by the spatial configuration, for example, through the general principles of **ease of reference** and **minimal effort**. Thus, a relatum is preferentially chosen that enables reference by simple, unmodified spatial terms, projective or other (e.g., distance-related or in-between relations).

The spatial relationship furthermore comes into play in the way the relationship to the prototypical spatial axis (from the relatum) is represented. With respect to descriptions of spatial relationships ("Where" questions) it has been noted that projective terms are associated with spatial templates, where a close relationship to the focal axis represents a typical relationship that can be expressed by unmodified linguistic descriptions; more deviant relationships require more complex descriptions (e.g., Vorwerg 2001). In the present discourse task, the only restriction seems to be that defined by a half plane. Otherwise, it is the presence of other (competing) objects, not the exact spatial relationship to the focal axis, which influences the employment of modifications in the linguistic descriptions. Thus, the more competing objects are present on the same half plane, the more precise the description will need to be, as required by the principle of **contrastivity**. If the goal object is the only one on a half plane, an unmodified term might be considered sufficient. This, again, reflects the principle of **minimal effort** that influences speakers' choices in spatial descriptions (just as well as other linguistic tasks).

Apart from the kind of **partner adaptation** involved in perspective choice, the participants in the HRI studies also adapted to the robot in other ways. They started out from their own ideas with respect to what the robot might be able to understand; subsequently, they based their linguistic and conceptual choices on their previous experience of success and failure, thereby adapting their spatial descriptions to the capabilities of the interaction partner. This led to a more restricted range of variety in speakers' descriptions (following their first success) as compared to the web studies, where no element of success or failure was involved. In this area, the main effect of mode comes into play: since the spoken mode involves additional problems of speech recognition, it is harder for users to update their knowledge about the robot's capabilities, and therefore harder to adapt the descriptions to the interaction partner's requirements.

One further interesting aspect which the two studies shed light on is dimensionality. In the web study, participants were confronted with a 2D scene, which should in a selection of the conditions be conceptualized as a 3D situation. It turned out that speakers did not have major difficulties with this imagination task. In fact, the only difference between 2D and 3D that was reflected in the linguistic representations concerned choice of axis: If the scenario was perceived as a scene in a picture (2D), then speakers chose the vertical axis for reference. In a 3D conceptualization, the frontal axis was employed, much in the same way as the vertical axis in 2D. The 3D concep-

tualizations, in turn, did not differ in fundamental respects from the real world scenarios employed in the HRI studies. Thus, it seems that even in a real world scenario, three-dimensional concepts do not influence the linguistic representation much; it is typically the horizontal plane that is crucial for spatial descriptions, in similar ways as the 2D (vertical-lateral) plane in pictures. Of course, this is influenced by the scenarios used in the present work because of the fact that the vertical dimension (in the real world scenarios) was never discriminative, i.e., all objects were situated on the same horizontal plane. But this seems to be a very typical situation, since scenarios in which the spatial position of a flying (or otherwise vertically offset) object needs to be defined using qualitative descriptions are rather exceptional. It would be interesting, of course, to address a scenario in which objects were positioned on three dimensions, especially within an indoor situation similar to the one used in the present work. To my knowledge this has not been done in the literature so far.¹

Altogether, the **speaker** perspective on the spatial task at hand seems to be the following. The speaker needs to determine which description is most suitable or least ambiguous with respect to the given goal object, also taking into account the interlocutor's capabilities. In some complex cases, this might also mean identifying a description that might be suitable for the discourse task at all, provided that the hearer is cooperative and accepts imprecise and fragmentary descriptions. The **hearer** perspective in the discourse is then to identify which object may be the best match for the given description, or, in some complex cases, to identify an object that could – with a generous interpretation – be described by the given description at all. If the assessments made by speaker and hearer do not match, communication fails, and negotiation will be needed to cooperatively solve the discourse task.

Because the present kind of discourse task has previously not been addressed in detail, many of these results are new with respect to their significance in spatial descriptions, especially with respect to the details discussed extensively in Tenbrink (2005c). Nevertheless, they are generally in line with previous results on basic processes involved in discourse. The principles concerning contrastivity and partner adaptation are already established in some detail in Herrmann & Deutsch (1976) with respect to other kinds of (non-spatial) object reference. The principles of minimal effort and further processes concerning adaptation to the interaction partner are addressed, for instance, in Clark & Wilkes-Gibbs (1986) and Clark (1996). Related processes of interactive alignment are spelled out in Pickering & Garrod (2004). Thus, the present work mainly spells out how these principles and processes combine in the application of projective terms.

Conclusion

I have summarized the results of two studies designed to trigger spontaneous production of projective terms, in order to gain insights about the ways in which they are applied in spontaneously produced language. As seems to be inevitable for any at-

¹ It would be interesting, however, to establish the relationship to air traffic control studies. To my knowledge, typically descriptions in air traffic situations rely predominantly on quantitative measures of height or angles, but this aspect could be examined more closely.

tempt to bridge the gap between natural discourse and laboratory findings, a number of objections could be raised: From a psycholinguistic point of view, both studies are not sufficiently controlled to yield (statistically) reliable results. But from a discourse analytic point of view, both settings may be considered to be too restricted and unnatural to allow for generalizations of the speakers' spontaneous choices in spatial reference. Nevertheless, the results of the present work show that not only are the speakers' choices and linguistic strategies within two completely different discourse settings comparable and consistent, but they can also be related directly to previous work, spelling out generalized hypotheses and findings gained in other settings and transferring them to new domains. Notably, each discourse setting has its own complications and features which need to be accounted for in the analysis.

Altogether, the present work points towards the value of using controlled settings but unconstrained language, as long as the findings can be analyzed in close relation to established results concerning both the setting and the specific research issues at hand. The rewards of such an approach are, on the one hand, the ability to verify results gained in different kinds of settings (e.g., psycholinguistic comprehension or constrained production tasks) and to spell them out for new application areas such as a broader range of discourse tasks, and on the other hand, the exploration of a particular discourse situation in order to generate specific hypotheses that can then be addressed directly in future research. In addition to that, for the purposes of technological progress, often already the results gained by exploration can be integrated directly into a robotic system under development, since the main aim in this field is typically not to identify underlying psychological processes but – pragmatically – to achieve an efficient system.

References

- Anderson, A. H., M. Bader, E. G. Bard, E. Boyle, G. Doherty, S. Garrod, S. Isard, J. Kowtko, J. McAllister, J. Miller, C. Sotillo, H. Thompson, and R. Weinert. 1991. The HCRC Map Task Corpus. *Language and Speech* 34(4), 351-366.
- Clark, H.H. 1996. *Using Language*. Cambridge, UK: Cambridge University Press.
- Clark, H. and Wilkes-Gibbs, D. 1986. Referring as a collaborative process. *Cognition*, 22:1–39.
- Dale, R. and E. Reiter. 1995. Computational interpretations of the Gricean maxims in the generation of referring expressions. *Cognitive Science*, 18:233-263.
- de Vega, M. and M.J. Rodrigo. 2005. The bicycle pedal is in front of the table. Why some objects do not fit into some spatial relations. In L. Carlson & E. van der Zee (eds.), *Functional features in language and space: Insights from perception, categorization and development*. Oxford: Oxford University Press, pp. 25-37.
- Fischer, K. and R. Moratz. 2001. From Communicative Strategies to Cognitive Modelling. *Workshop Epigenetic Robotics*, Lund.
- 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.
- Gorniak, P. and D. Roy. 2004. Grounded Semantic Composition for Visual Scenes. *Journal of Artificial Intelligence Research*, Volume 21, pages 429-470.
- Herrmann, T. and W. Deutsch. 1976. *Psychologie der Objektbenennung*. Bern u.a.: Hans Huber Verlag.
- Montello, D. R. 1993. Scale and multiple psychologies of space. In Frank, A. and Campari, I., eds, *Spatial Information Theory: A theoretical basis for GIS*, pp. 312–321, Berlin, Heidelberg, New York. Springer.

- Moratz, R. and T. Tenbrink. 2003. Instruction modes for joint spatial reference between naive users and a mobile robot. *Proceedings of the IEEE International Conference on Robotics, Intelligent Systems and Signal Processing (RISSP 2003)*, Special Session on New Methods in Human Robot Interaction, October 8-13, 2003, Changsha, Hunan, China.
- Moratz, R. and T. Tenbrink. 2006. Spatial reference in linguistic human-robot interaction: Iterative, empirically supported development of a model of projective relations. *Spatial Cognition and Computation, Volume 6, Issue 1*, pp. 63-106.
- Pickering, M. and Garrod, S. 2004. Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27, 169-226.
- Schober, M.F., and Brennan, S.E. 2003. Processes of interactive spoken discourse: The role of the partner. In A.C. Graesser, M.A. Gernsbacher, & S.R. Goldman (eds.), *Handbook of Discourse Processes* (pp. 123-164). Mahwah, NJ: Lawrence Erlbaum Associates.
- Sperber, Dan and Ira A. Noveck. 2006. Introduction. In I.A. Noveck & D. Sperber, *Experimental Pragmatics*. Palgrave Macmillan.
- Tenbrink, T. 2005a. Identifying objects in English and German: A contrastive linguistic analysis of spatial reference. *Proc. Workshop on Spatial Language and Dialogue (5th Workshop on Language and Space)*, October 23-25, 2005, Delmenhorst, Germany.
- Tenbrink, T. 2005b. Identifying objects on the basis of spatial contrast: An empirical study. In Freksa, Christian, Markus Knauff, Bernd Krieg-Brückner, Bernhard Nebel, and Thomas Barkowsky (eds.), *Spatial Cognition IV: Reasoning, Action, Interaction*. International Conference Spatial Cognition 2004, Frauenchiemsee, Germany, October 2004, Proceedings. Berlin, Heidelberg: Springer, pp. 124-146.
- Tenbrink, T. 2005c. Localising objects and events: Discoursal applicability conditions for spatiotemporal expressions in English and German. *Dissertation*, University of Bremen.
- Tenbrink, T. 2005d. Semantics and application of spatial dimensional terms in English and German. *SFB/TR8 Report No. 004-03/2005*, Universität Bremen / Universität Freiburg.
- van Deemter, K. and R. Kibble. 2002. *Information Sharing: Reference and Presupposition in Language Generation and Interpretation*. Stanford: CSLI Publications.
- van der Sluis, I. and E. Kraemer. 2000. Generating Referring Expressions in a Multimodal Context: An empirically motivated approach. In: W. Daelemans et al. (eds.), *Computational Linguistics in the Netherlands 2000, Selected Papers from the 11th CLIN Meeting*. Rodopi, Amsterdam.
- Vorwerg, C. 2001. *Raumrelationen in Wahrnehmung und Sprache: Kategorisierungsprozesse bei der Benennung visueller Richtungsrelationen*. Wiesbaden: DUV.