

Deep Reasoning in Clarification Dialogues with Mobile Robots

Cui Jian¹ and Desislava Zhekova² and Hui Shi³ and John Bateman⁴

Abstract. This paper reports our work on qualitative reasoning based clarification dialogues in human-robot interaction on spatial navigation. To interpret humans’ route instructions, a qualitative spatial model is introduced which represents the robot’s beliefs in the application domain. Based on the qualitative spatial model, three tool-supported reasoning strategies are discussed which enable the robot to generate dialogues with differing degrees of clarification if knowledge mismatches or under-specifications in route instructions are detected. The influence of the reasoning strategies on human-robot dialogues is evaluated with an empirical study.

1 Motivation

While instructing a mobile robot to navigate in a partially known environment, humans are likely to make some *knowledge-based mistakes*, since describing a route is a high-level cognitive process and involves the assessment of complex environment information—such as different spatial frames of reference, localization of spatial objects, and spatial relations between these objects [13].

Several research efforts on natural language communication between humans and intelligent systems have been reported in the literature. Some have focused on issues such as corpora based primitive functions (e.g., [12, 4]), task-oriented slot-filling based planning assistants (e.g., [2]) or task tree based tutoring systems (e.g., [6]), while others have applied logical formalization and reasoning in dialogue systems (e.g., [16, 9]). In contrast to these directions, we focus here on the spatial application domain and address the particular challenge of developing a conceptual knowledge model that is capable both of maintaining the spatial knowledge of the environment and of providing a semantic framework for interpreting and reasoning about spatial information to support human-robot dialogues.

Conceptual knowledge representations are already used in some systems to enable human-robot interaction. For example, Zender *et al.* [22] present an approach to create conceptual spatial representations of indoor environments with different layers of map abstractions to support human-robot dialogues. Here we apply instead mathematically well-founded qualitative spatial calculi and models (cf. [1, 10, 23, 7]) since these provide not only the semantic interpretation of humans’ spatial knowledge but also mechanisms for reasoning with this knowledge. Although qualitative spatial models have been used for representing and reasoning about spatial relations before, their application as a belief model to assist a robot to communicate naturally with a human has received little attention.

This paper therefore focuses on the development of a robot’s spatial belief model taking qualitative spatial calculi as its foundation. We discuss various tool-supported reasoning strategies which enable the robot to validate route instructions according to its spatial knowledge and create spatial relations that indicate the mismatches or under-specifications in those instructions. Clarification dialogues are then generated according to the relations created. We explore how these more informative clarification dialogues can be provided using qualitative reasoning and evaluate their contribution to effective human-robot interaction systems.

The paper is structured as follows: Section 2 defines the qualitative spatial beliefs model, QSBM, and a set of update rules for interpreting common route instructions based on QSBM. Section 3 introduces three different reasoning strategies and their implementation in the system *SimSpace*. Section 4 then describes an evaluation study and several evaluation results are presented in Section 5. We discuss our approach and results more generally in Section 6, before concluding in Section 7.

2 Modelling a Robot’s Spatial Beliefs

Unlike dialogue history or general conversational information in dialogue systems, domain knowledge is application dependent; this requires that a dialogue system be able to talk with the user about domain specific issues. These issues therefore require appropriate representation in their own right. Modern intelligent mobile robots can navigate autonomously in various environments using sensor techniques; but for people to communicate with such robots for joint spatial tasks, an intermediate knowledge representation is necessary so that they can understand each other. The qualitative spatial beliefs model presented in this section is a representation for facilitating such communication.

2.1 Route Graph and Double-Cross Calculus

One proposal for a common knowledge base for humans or mobile agents involved in navigation is that of *route graphs* [20]. Such graphs capture humans’ topological knowledge on the qualitative level while acting in space. Route graphs are a special class of graphs. A *node* of a route graph, called a *place*, has a particular position and its own “reference system” for defining directions at that place. An *edge* of a route graph, called a *route segment*, is directed from a source place to a target place and always has three attributes: an *entry*, a *course* and an *exit*.

The Double-Cross calculus (DCC) was introduced by Freksa [11, 23] for qualitative spatial representation and reasoning using orientation information. Combining the front/back and the left/right

¹ University of Bremen, Germany, email:ken@informatik.uni-bremen.de

² University of Bremen, Germany, email:zhekova@uni-bremen.de

³ DFKI Bremen, Germany, email:hui.shi@dfki.de

⁴ University of Bremen, Germany, email:bateman@uni-bremen.de

dichotomy, the DCC distinguishes 15 meaningful qualitative orientation relations (or DCC relations), such as “front”, “right front”, “right”, etc.

The Conceptual Route Graph (CRG) of Shi and Krieg-Brückner [14] then combines the structure of conventional route graphs and the Double-Cross calculus. The entry and exit of a route segment is further defined by an orientation between the route segment and the reference frame at its source or target place. Additionally, a set of DCC relations are used to describe the orientation relations between route segments and places. A *Route* of a CRG is then a sequence of connected route segments. Thus, CRGs can be seen as route graphs with only qualitative information, i.e. in the present case, the DCC relations.

2.2 QSBM: Qualitative Spatial Beliefs Model

We define a QSBM as a pair of a conceptual route graph and a route segment representing the current position of the robot. We denote this by $\langle \text{crg}, \text{pos} \rangle$, where the robot is currently located at the entry place of the route segment pos and oriented in the direction of its exit place. A conceptual route graph is simply represented by a tuple of four elements $(\mathcal{M}, \mathcal{P}, \mathcal{V}, \mathcal{R})$, where $\mathcal{M}$ is a set of landmarks in the environment, $\mathcal{P}$ a set of topological places, $\mathcal{V}$ a set of vectors from a source place to a target place, and $\mathcal{R}$ a set of orientation relations.

Let p be a place and $\overline{xy}$ a vector from place x to place y . The typical orientations of p with respect to $\overline{xy}$ are then: p is right of $\overline{xy}$ (written as $\langle \overline{xy}, \text{Right}, p \rangle$), p is on $\overline{xy}$ (written as $\langle \overline{xy}, \text{On}, p \rangle$), or p is in front of $\overline{xy}$ (written as $\langle \overline{xy}, \text{Front}, p \rangle$), etc. according to the 15 orientation fields distinguished by the Double-Cross calculus.

In the following discussion we focus on our application scenario, in which we presume that the robot knows essential spatial arrangements (i.e., significant places, connections and orientation relations) of the environment represented by the QSBM.

2.3 Update Rules

Various empirical studies (cf. [12, 4, 15]) confirm that route instructions given by a human to a mobile robot resemble how people instruct other people to navigate (cf. [18, 8]). A common human route description accordingly consists of a sequence of route instructions concerning egocentric or landmark-based reorientations, according to which the QSBM is then progressively updated. In this section we will discuss QSBM update rules for three classes of route instructions used as examples throughout this paper. Each rule has a name, a set of pre-conditions and an effect part. The symbols and operators used in the definitions of update rules are explained where they are first used; in addition we employ the standard operators $\exists$ for the existential quantifier, $\neg$ for its negation, and $\in$ for the element test operator on sets.

Reorientation is typically expressed by directional instructions like “turn left/right” and “turn around”, which may change the orientation of the robot at the current position. The pre-condition for applying this rule is that the robot should find a place in its belief model that satisfies the desired relation and the effect is that it faces that place after the turn operation. Suppose the robot is currently at the place p_0 and faces the place p_1 , then the pre-condition is stated as: there exists a place p such that the orientation of p with respect to the route segment $\overline{p_0 p_1}$ (or the current position) is the given direction d to turn, i.e., $\langle \overline{p_0 p_1}, d, p \rangle$. Formally, we use the rule *Reorientation* to specify the turn-operation, and assume that $\langle (\mathcal{M}, \mathcal{P}, \mathcal{V}, \mathcal{R}), \text{pos} \rangle$ is the belief model of the robot with the current position pos .

RULE: Reorientation

PRE: $\text{pos} = \overline{p_0 p_1}, \exists p \in \mathcal{P}. \langle \overline{p_0 p_1}, d, p \rangle$
 EFF: $\text{pos} = \overline{p_0 p}$

Moving through instructions, such as “go through the door”, usually contain a landmark which should be in front of the robot before, and behind the robot after, the move action. Thus, an important pre-condition here is to find a route segment (say $\overline{p_2 p_3}$) in front of the current position $\overline{p_0 p_1}$, such that the landmark is on $\overline{p_1 p_2}$. $\overline{p_2 p_3}$ is then the new robot position, as specified in the rule *MoveThrough*.

RULE: MoveThrough

PRE: $\text{pos} = \overline{p_0 p_1},$
 $\exists p \in \mathcal{P}, \overline{p_2 p_3} \in \mathcal{V}. \text{at}(l, p)$
 $\wedge \langle \overline{p_1 p_2}, \text{On}, p \rangle \wedge \langle \overline{p_0 p_1}, \text{Front}, p_2 \rangle \wedge \langle \overline{p_1 p_2}, \text{Front}, p_3 \rangle$
 EFF: $\text{pos} = \overline{p_2 p_3}$

In the rule above l is the landmark referred to in the instruction. The logical conjunction-operator $\wedge$ is used to define the second pre-condition. Moreover, the relation *at* associates a landmark with its location.

Passing classifies route instructions containing path external spatial descriptions, such as “pass the copy room on the right” or just “pass the copy room” without direction information. Again, as a pre-condition the robot should identify the landmark given in the instruction and check whether there is a route such that the landmark is before it at the beginning and behind it at the end of the motion (see the second pre-condition in the following rule). If the direction of passing the landmark is presented in the instruction, certain orientation relations should be satisfied as well. The rule *PassRight* specifies the update strategy for the case in which a landmark should be passed on the right.

RULE: PassRight

PRE: $\text{pos} = \overline{p_0 p_1},$
 $\exists p \in \mathcal{P}, \overline{p_2 p_3} \in \mathcal{V}. \text{at}(l, p)$
 $\wedge \langle \overline{p_0 p_1}, \text{RightFront}, p \rangle \wedge \langle \overline{p_2 p_3}, \text{RightBack}, p \rangle$
 $\wedge \langle \overline{p_0 p_1}, \text{Front}, p_2 \rangle \wedge \langle \overline{p_1 p_2}, \text{Front}, p_3 \rangle$
 EFF: $\text{pos} = \overline{p_2 p_3},$
 if $\neg \overline{p_4 p_2} \in \mathcal{V}. \langle \overline{p_4 p_2}, \text{RightBack}, p \rangle \wedge \langle \overline{p_0 p_1}, \text{Front}, p_4 \rangle$

Consequently, the current position of the robot will be updated using the shortest route that satisfies the condition.

3 Qualitative Spatial Reasoning about Route Instructions

As stated in Section 1, humans’ route instructions may contain knowledge-based mistakes or incomplete information. Take the sample instance of an environment in Fig. 1, where “the copy room” (at place g) and “the mailbox room” (at h) are located left front of the robot whose current position is denoted as $\overline{ab}$. The route instruction “pass by the copy room on the right” causes an orientation mismatch, while “pass the copy room” contains no information about the spatial orientation of the copy room. In this section we introduce three reasoning strategies, based on QSBM, to perform high level interpretation of such route instructions with respect to the introduced update rules (cf. Section 2.3). We also discuss their corresponding implementation in the system *SimSpace*.

3.1 Three Reasoning Strategies

In order to perform the interpretation of a sequence of route instructions, the environment state is checked and then updated based on QSBM. According to how such processes are executed, three reasoning strategies are distinguished.




Figure 1. A sample conceptual environment

3.1.1 Shallow Reasoning

The *shallow reasoning* strategy is the simplest one of the three reasoning strategies we discuss. It just checks the condition for interpreting a route instruction according to the environment state. If the condition can be satisfied by the environment, the instruction is successfully interpreted, otherwise the route instruction is rejected.

We take the environment in Fig. 1 again and interpret the route instruction “pass the copy room on the right”. A place has to be found at which the copy room is located, and whose relation with the current position of the robot $\overline{ab}$ is *Right*. Such a place cannot be found in this situation and so the system rejects the instruction.

3.1.2 Deep Reasoning

In contrast, the *deep reasoning* strategy does not simply reject a route instruction, but tries to generate correct spatial relations in contrast to the inconsistent ones contained in the route instruction, i.e., deep reasoning provides relevant spatial information about the specific situation.

Taking the above “pass the copy room on the right” example again, the deep reasoning strategy detects the copy room at the place g and its spatial relation with the robot’s current position $\overline{ab}$ as left front, which is known from the DCC definitions to be inconsistent with the orientation relation contained in the route instruction.

The deep reasoning strategy then enables the reasoning system to deliver spatial relations that describe specific spatial configurations, and which can be used to explain why a given route instruction cannot be interpreted, assisting people to respond with a more appropriate spatial configuration.

3.1.3 Deep Reasoning with Backtracking

Providing a sequence of route instructions to navigate a robot to a certain place is complex: people not only need to locate the current and the destination position correctly – they also have to work in a dynamic mental situation, in which appropriate route instructions need to be constructed and connected with the relevant route and the imagined current position has to be updated after each instruction is executed mentally. If people make a wrong mental rotation [21] by using a wrong route instruction, and they often do [13], the remaining route instructions will not lead to the desired destination; they might even be uninterpretable since they do not match the subsequent spatial situation.

Consider the route instruction “take the first junction on the left and then drive straight to the mailbox room” in Fig. 1. No interpretation can be achieved in this situation because, after the robot takes the first junction to the left, it cannot find any mailbox room from the current perspective. However, by taking one step backwards, instead of “take the *first* junction on the left”, “take the *second* junction

on the left and then drive straight to the mailbox room” satisfies the constraints of the situation perfectly.

Therefore, when a route instruction is uninterpretable in a certain situation, instead of simply providing a reason as in *deep reasoning*, our third strategy, *deep reasoning with backtracking*, tries to locate this potentially incorrect instruction, runs the corresponding forward checking of the subsequent route instructions with all possible corrections, and finally achieves a successful interpretation if it exists.

3.2 SimSpace: An Implementation

The qualitative spatial belief model QSBM and its accompanying reasoning strategies has been implemented in the tool *SimSpace*. This is now able to interpret most commonly used route instructions including those discussed in this paper, and to generate *suggest* or *inform* responses if knowledge mismatches or underspecifications are detected.

Given a route instruction parsed in a pre-defined semantic form, *SimSpace* interprets it by automatically selecting an applicable update rule and instantiating its pre-conditions. Taking the sample route instruction “pass the copy room on the right” as an example, *SimSpace* selects the update rule *PassRight* to interpret it. Assuming $\overline{ab}$ to be the current position, the second pre-condition is instantiated to

$$\begin{aligned} &\exists p \in \mathcal{P}, \overline{p_2 p_3} \in \mathcal{V}. at(\text{CopyRoom}, p) \\ &\wedge \langle \overline{ab}, \text{RightFront}, p \rangle \wedge \langle \overline{p_2 p_3}, \text{RightBack}, p \rangle \end{aligned}$$

Our adoption of a standard qualitative calculus then allows us to employ newly emerging generic tools for qualitative reasoning. Thus, using the qualitative spatial reasoner *SparQ* [19], the instantiated pre-conditions are checked against the current state of the QSBM and one of our pre-specified strategies. In the case when the shallow reasoning strategy is used, the instruction will be simply rejected by the robot with the response “Not possible”. However, using the deep reasoning strategy to interpret the same instruction, we get the result:

$$at(\text{CopyRoom}, g), \langle \overline{ab}, \text{LeftFront}, g \rangle$$

i.e., the copy room is located at g and on the left side of the current position. Thus, the orientation relation $\langle \overline{ab}, \text{RightFront}, p \rangle$ in the pre-condition cannot be satisfied. Consequently, *SimSpace* translates these results into a corresponding expression in the sentence planning language *SPL* in order to allow generation of a corresponding natural language response using the *KPML* natural language generator [3]. In this case, the response: “I cannot pass the copy room on the right, but I can pass it on the left.” is generated.

Finally, by adding the backtracking strategy, every state of QSBM and its interpreted instruction are recorded in a transition history before interpretation proceeds. If any checking/updating process fails, *SimSpace* reloads the previous states and calculates a possible alternation/correction with respect to the interpreted instruction and, if possible, proceeds with the interpretation of the found alternation/correction. Therefore, the response to the instruction “take the first junction on the left and then drive straight to the mailbox room” is then “if I take the second junction on the left, I can drive straight to the mailbox room”.

4 An Empirical Study

In order to explore how human-robot dialogues concerning navigation tasks are influenced by our three reasoning strategies, an empirical study was carried out as follows.

4.1 Participants

6 scientific researchers and 15 students, i.e., a total of 21 participants (11 female, 10 male) took part in the study voluntarily.

4.2 Stimuli and Apparatus

Three maps of common indoor environments, called the navigation maps, were coupled with the three reasoning strategies and used throughout the study by each participant for the navigation tasks. Each map contained 20 locations (7 named, 13 unnamed) and was made with similar spatial configurations with the same level of complexity (cf. [5]); in addition, in order to help the participants to memorize the maps, all maps were designed with a common object shape layout according to the principle of imagery mnemonics (cf. [17]); cf. “The Chimney House Clinic” in Fig. 2. Moreover, to retain the start configuration, the position and orientation of a simulated robot was prescribed on each map and this remained the same for all participants.



Figure 2. The Chimney House Clinic

The simulated dialogue system was a networked software application that connected two computers: one computer, called the *navigation assistant*, constantly displayed the current system response and a list of the given route instructions in natural language to the participants; the other, called the *brain system*, was controlled by a human operator who entered the route instructions given by the participant. The *SimSpace* tool (Section 3.2) is the key component of the brain system, checking the instructions using a pre-specified reasoning strategy according to the robot's QSBM and generating clarification dialogues if necessary. As a result, the whole test run was simulated as if the participant communicates with the system in natural language directly, but removing possible distractors that might have been introduced by speech recognition or parsing errors.

An evaluation questionnaire concerning the participant's memorization of the map used and his/her feeling about the system responses was then completed.

4.3 Procedure

Each test was divided into two phases: learning and testing.

4.3.1 Learning Phase

Each participant was given a general introduction to the test procedure, including the navigation tasks, the spatial configuration of the maps, commonly used route instructions (cf. [15]) and the ways to interact with the navigation assistant. A sample map was also presented to the participant, who was then asked to accomplish several prescribed sample navigation tasks. The test would only move on to the next phase when the participant had acquired the information necessary.

4.3.2 Testing Phase

In this phase each participant had to go through three test runs, which were coupled with the three navigation maps and the three reasoning strategies. The sequence of the maps and the combination between each map and each reasoning strategy remained unchanged throughout the study. Each test run consisted of the following three steps:

- *Memorization.* The participant was asked to memorize a given map of a common indoor environment within exactly one minute; then the map was taken away.
- *Navigation.* The participant was free to communicate with the navigation assistant by giving oral route instructions to navigate the mobile robot to three different places on the memorized map. Each task was only finished when the destination was reached or the participant gave up trying.
- *Evaluation.* At the end of each test run, the participant was asked to fill in the evaluation questionnaire.

5 Results and Analysis

Regarding the dialogue that was formed in the context of the second step of the tests—the *Navigation* step—several interesting points emerged.

In respect to the first and most simplistic reasoning strategy (shallow reasoning; cf. Section 3.1.1), we observed repetition of the mistakes the participants made in their instructions, confusion, and failure to reach the goal as the three most noticeable problems. The repetitiveness of mistakes was mainly caused by the uninformative nature of the system answers (*OK* or *Not possible*). Consider the following dialogue turns:

User:	Drive through the door.
System:	OK.
User:	Turn left.
System:	OK.
User:	Drive to the end of the corridor.
System:	OK.
User:	Turn right.
System:	OK.
User:	Take the second turning on the right.
System:	Not possible.

Since from the given dialogue with the system the user can conclude that only the last instruction was wrong, he/she makes an attempt to change only this one. Consequently in the next trial one of the possible replacement utterances could be the following:

User:	Take the third turning on the right.
System:	Not possible.

The latter is again wrong, but the user does not receive more information about what exactly is wrong with it. Therefore the user keeps

the same course of action and changes again and again the last instruction.

The lack of informativeness in the system answers led as well to another negative result: the confusion and helplessness that accumulated with each negative system answer. Let us consider the following situation:

User: Drive through the door.

System: Not possible.

In this case the user correctly navigated to the dining room. The only mistake was that the robot was not oriented in the direction of the goal (see the update rule *MoveThrough* in Section 2.3). However, receiving an answer “*Not possible*” to the instruction “*Drive through the door*” confuses the user (i.e. what the user might think is that if it is not possible to enter the dining room then the position of that room must be somewhere else). Thus, the navigator tries to change the correctly remembered position of the dining room and find another position for it in the mental representation of the map he/she has. In other words, the user creates a new spatial mismatch as a result of the system answer. The further negative answers of the system, which lack any reasoning about why it is not possible to drive through the door at that point, only make the user give up trying to reach the goal. Consequently only 42 of the 63 goals altogether were reached by the participants with the help of the shallow reasoning strategy.

Only giving instructions and getting “*OK*” or “*Not Possible*” as answers can hardly be considered helpful dialogue. On the other hand, additionally giving a reason for why it is not possible to perform a certain action certainly can. Reasoning does not only improve on the naturalness of the dialogue, but on its usefulness as well. We consider two very simple instructions in the first and second reasoning strategies in Table 1. Naturally, after receiving the answers to the instructions in Strategy #1 (shallow reasoning), the user does not know what went wrong and what instruction would result in a successful completion of the task. In Strategy #2 (deep reasoning) however, the user receives enough information in order to give an instruction that could lead either to the completion of the task or to a correction of the previous false instruction. We consider the latter to be an important factor that participates considerably to the increase in the number of trials in which goals were successfully reached: i.e., 48/63.

Str.	Answer
#1	Not possible.
#2	We can not drive until the lab, because it is behind us.

Table 1. Example system answers from Strategy #1/#2 to the instruction *Drive until the lab*.

A further improvement of the dialogue capabilities is provided in the third reasoning strategy—Strategy #3 (deep reasoning with backtracking); this is the possibility for correcting false instructions by giving a suggestion for an action that is closest to the given command and also possible to perform (see Section 3.1.3). Examples of such dialogue turns are given in Table 2. In natural conversations humans are normally able to give a reason if they can not perform a certain action, but it is not always the case that they can suggest what could be done to make corrections due to the lack of such knowledge. Thus, we consider that the deep reasoning strategy with backtracking moves us closer to a cognitive and helpful natural language human-human dialogue. As shown in Table 3, 58 goals from altogether 63 were found in comparison to the 42 goals out of 63 with the simplest reasoning strategy. This supports our assessment of Strategy #3 as being more helpful to participants.

Str.	Answer
#1	Not possible.
#3	If I take the first junction to the left, I can’t drive until the Communication department on the right, but if I take the second junction to the left, I could drive until the Communication department on the right.

Table 2. Example system answer from Strategy #1/#3 to the instruction *Take the first junction to the left and then drive to the Communication Department to the right*.

In the second row of Table 3 are figures representing the average number of instructions in the cases in which the participants reached the goal. The fact that this number is increasing for the second and third reasoning strategy demonstrates that the users actually had longer dialogues with the system to reach the goals and did not rely only on their own memory but on the information provided in the system answers as well.

Last but not least, we can pay attention to the last row of Table 3 where we show the satisfaction degree of the participants indicating how content they were with the system answers for each reasoning strategy. According to these questionnaire results, the shallow parsing strategy is ranked as least sufficient with a result almost 15% below the next reasoning approach. Deep reasoning with backtracking received only about 2% higher result than deep reasoning. This according to our observations is an effect brought about by the increasing complexity of the system responses when backtracking is added. Nevertheless, both latter strategies appeared to be appreciated well by the participants.

The results of this experiment serve to demonstrate the influence of clarification dialogues with mobile robots as well as to point out the direction in which the system should be further developed. They accordingly provide a base for more exhaustive evaluation that has been planned for the further steps of the development.

	Str. #1	Str. #2	Str. #3
Reached Goals	66.67%	76.19%	92.06%
Average Instructions per Reached Goal	10.17	11.15	13.67
Satisfaction Degree	58.2%	73.0%	75.3%

Table 3. Summary of the experimental results.

6 Discussion

Spatial mismatches can often be observed in human route instructions in everyday life, but they usually do not pose a great problem in human-human communication since humans are able to easily spot and clarify the mismatch at hand. In human-robot interactions during navigation tasks, however, such spatial mismatches can cause difficulties and complications as well as predisposing users negatively to the system.

In our work we have shown that deep reasoning with and without backtracking in clarification dialogues with mobile robots can help spot, exemplify and reduce to a great extent the spatial mismatches in human-robot dialogue. Moreover, once the mismatch is identified, an informative and helpful correction and suggestion increases the usefulness of the dialogue itself. This makes the latter more constructive, valuable and helpful—characteristics still problematic in the area of human-robot dialogue systems at large.

Furthermore, we have also treated some new specific challenges for situated dialogue in the current work. The issue is not only that more informative feedback helps users complete a larger number of direction-giving tasks, which is indeed not surprising, but precisely how that needs to be done with what kind of additional information. It is crucial to explore this precise point of contact if effective dialogue systems for human-robot are to be built: we need to know more about just what information is needed where. This is explored directly in the experimental setup pursued.

Thus, such reasoning-based informative dialogues should be considered a highly desirable design feature for situated and general dialogue systems that address the need of improving human-robot dialogue so as to resemble human-human interaction more closely.

7 Conclusion and Future Work

This paper has reported work that integrates several interests of Artificial Intelligence. Concretely we treat the following aspects in depth: the management and formalization of, and reasoning with, domain knowledge. All three represent essential components for human-robot dialogue systems. Specifically, we presented three reasoning strategies and discussed their influences on clarification dialogues with mobile robots. The major contributions of the current work are twofold: the development of the robot's qualitative spatial belief model (QSBM), and the generation of the robot's responses using qualitative reasoning. Furthermore, a preliminary empirical study confirmed that qualitative spatial reasoning mechanisms, in our case the deep reasoning and the deep reasoning with backtracking strategies, provide useful information for the robot to generate more natural and informative dialogues. This therefore encourages further experiments based on these results.

The integration of the qualitative spatial belief model, including the reasoning mechanisms and their implementation in *SimSpace*, into a natural language dialogue system is still ongoing work. We are also planning to add more refined qualitative calculi and reasoning abilities to enrich both the descriptions of the application domains and the sophistication of the system's spatial responses. Moreover, the research on the close interaction between knowledge management and dialogue modelling/control, e.g., the application of our reasoning methods to catch misunderstandings of the system, will also be pursued further.

ACKNOWLEDGEMENTS

We gratefully acknowledge the support of the Deutsche Forschungsgemeinschaft (DFG) through the Collaborative Research Center SFB/TR 8 Spatial Cognition - Subprojects I3-SharC and I5-DiaSpace, and Department of Safe & Secure Cognitive Systems at German Research Center for Artificial Intelligence (DFKI) Bremen.

REFERENCES

- [1] James F. Allen, 'Maintaining knowledge about temporal intervals', *CACM*, **26**(11), 832–843, (1983).
- [2] James F. Allen, Lenhart K. Schubert, George Ferguson, Peter Heeman, Chung Hee Hwang, Tsuneaki Kato, Marc Light, Nathaniel G. Martin, Bradford W. Miller, Massimo Poesio, and David R. Traum, 'The trains project: A case study in building a conversational planning agent', *Journal of Experimental and Theoretical AI*, **7**, 7–48, (1994).
- [3] John A. Bateman, 'Enabling technology for multilingual natural language generation: the KPML development environment', *Journal of Natural Language Engineering*, **3**(1), 15–55, (1997).
- [4] Guido Bugmann, Ewen Klein, Stanislaw Lauria, and Theodor Kyriacou, 'Corpus-based robotics: A route instruction example', in *Proceedings of IAS-8*, (2004).
- [5] John Cambell, *Map Use and Analysis*, Columbus(OH): McGraw Hill, 2001.
- [6] Brady Clark, Oliver Lemon, Alexander Gruenstein, Elizabeth Owen Bratt, John Fry, Stanley Peters, Heather Pon-Barry, Karl Schultz, Zack Thomsen-Gray, and Pucktada Treeratpituk, *Advances in Natural Multimodal Dialogue Systems*, chapter 13: A General Purpose Architecture for Intelligent Tutoring Systems, 287–305, Springer Netherlands, 2005.
- [7] Anthony G. Cohn, Brandon Bennett, John Gooday, and Nicholas Mark Gotts, 'Qualitative spatial representation and reasoning with the region connection calculus', *GeoInformatica*, **1**(3), 275–316, (1997).
- [8] Michel Denis, 'The description of routes: A cognitive approach to the production of spatial discourse', *Cahiers de Psychologie Cognitive*, **16**, 409–458, (1997).
- [9] Debora Field and Allan Ramsay, 'Deep-reasoning-centered dialogue', in *The 11th European Workshop on Natural Language Generation*, pp. 131–138, Morristown, NJ, USA, (2007). Association for Computational Linguistics.
- [10] Andrew U. Frank, 'Qualitative spatial reasoning with cardinal directions', in *Proceedings of the Seventh Austrian Conference on Artificial Intelligence*, Springer, (1991).
- [11] Christian Freksa, 'Qualitative spatial reasoning', in *Cognitive and Linguistic Aspects of Geographic Space*, eds., D. M. Mark and A. U. Frank, Kluwer, (1991).
- [12] Stanislaw Lauria, Theodor Kyriacou, Guido. Bugmann, Johan Bos, and Ewan Klein, 'Converting natural language route instructions into robot executable procedures', in *Proceedings of the 2002 IEEE International Workshop on Human and Robot Interactive Communication*, pp. 223–228, (2002).
- [13] James Reason, *Human Error*, Cambridge University Press, 1990.
- [14] Hui Shi and Bernd Krieg-Brückner, 'Modelling Route Instructions for Robust Human-Robot Interaction on Navigation Tasks', *International Journal of Software and Informatics*, **2**(1), 33–60, (2008).
- [15] Hui Shi and Thora Tenbrink, 'Telling Rolland where to go: HRI dialogues on route navigation', in *Spatial Language and Dialogue*, eds., K. Coventry, T. Tenbrink, and J. Bateman, Cambridge University Press, (2009).
- [16] Ronnie W. Smith and D. Richard Hipp, *Spoken Natural Language Dialog Systems: A Practical Approach*, Oxford University Press, 1994.
- [17] John T.E. Richardson, *Cognitive Psychology - A Modular Course - Imagery*, chapter 5, 103–136, Psychology Press, Hove, UK, 1999.
- [18] Barbara Tversky and Paul U. Lee, 'How space structures language', in *Spatial Cognition: An interdisciplinary Approach to Representation and Processing of Spatial Knowledge*, eds., C. Freksa, C. Habel, and K.F. Wender, volume 1404 of *Lecture Notes in Artificial Intelligence*, pp. 157–175. Springer-Verlag, (1998).
- [19] Jan Oliver Wallgrün, Lutz Frommberger, Diedrich Wolter, Frank Dylla, and Christian Freksa, 'A toolbox for qualitative spatial representation and reasoning', in *Spatial Cognition V: Reasoning, Action, Interaction: International Conference Spatial Cognition 2006*, eds., T. Barkowsky, M. Knauff, G. Ligozat, and D. Montello, volume 4387 of *Lecture Notes in Computer Science*, pp. 39–58, (2007).
- [20] Steffen Werner, Bernd. Krieg-Brückner, and Theo Hermann, 'Modelling navigational knowledge by route graphs', in *Spatial Cognition II: Integrating Abstract Theories, Empirical Studies, Formal Methods, and Practical Applications*, volume 1849 of *Lecture Notes in Artificial Intelligence*, pp. 259–316. Springer-Verlag, (2000).
- [21] Christopher D. Wickens, 'Frames of reference for navigation', in *Attention and Performance XVII: Cognitive Regulation of Performance: Interaction Of Theory and Application*, eds., Daniel Gopher and Asher Koriati, 112–144, MIT Press, Cambridge, MA, (1999).
- [22] Hendrik Zender, O. Martinez Mozos, Patric Jensfelt, G.-J. M. Kruijff, and Wolfram Burgard, 'Conceptual spatial representations for indoor mobile robots', *Robotics and Autonomous Systems*, **56**, (2008).
- [23] Kai Zimmermann and Christian Freksa, 'Qualitative spatial reasoning using orientation, distance, and path knowledge', *Applied Intelligence*, **6**, 49–58, (1996).