

# Situated Dialogue and Spatial Organization: What, Where... and Why?

Geert-Jan M. Kruijff<sup>1</sup>; Hendrik Zender<sup>1</sup>; Patric Jensfelt<sup>2</sup> & Henrik I. Christensen<sup>2</sup>

<sup>1</sup>Language Technology Lab, German Research Center for Artificial Intelligence (DFKI GmbH), Saarbrücken, Germany

<sup>2</sup>Centre for Autonomous Systems, Royal Institute of Technology (KTH), Stockholm, Sweden  
gj@dfki.de

**Abstract:** *The paper presents an HRI architecture for human-augmented mapping, which has been implemented and tested on an autonomous mobile robotic platform. Through interaction with a human, the robot can augment its autonomously acquired metric map with qualitative information about locations and objects in the environment. The system implements various interaction strategies observed in independently performed Wizard-of-Oz studies. The paper discusses an ontology-based approach to multi-layered conceptual spatial mapping that provides a common ground for human-robot dialogue. This is achieved by combining acquired knowledge with innate conceptual commonsense knowledge in order to infer new knowledge. The architecture bridges the gap between the rich semantic representations of the meaning expressed by verbal utterances on the one hand and the robot's internal sensor-based world representation on the other. It is thus possible to establish references to spatial areas in a situated dialogue between a human and a robot about their environment. The resulting conceptual descriptions represent qualitative knowledge about locations in the environment that can serve as a basis for achieving a notion of situational awareness.*

**Keywords:** *Human-Robot Interaction, Conceptual Spatial Mapping, Situated Dialogue*

## 1. Introduction

More and more robots find their way into environments where their primary purpose is to interact with humans to help and solve a variety of service-oriented tasks. Particularly if such a service robot is mobile, it needs to have an understanding of the spatial and functional properties of the environment in which it operates. The problem we address is how a robot can acquire an understanding of the environment so that it can autonomously operate in it, and communicate about it with a human. We present an architecture that provides the robot with this ability through a combination of human-robot interaction and autonomous mapping techniques. It captures various functions that independently performed Wizard-of-Oz studies have observed to be necessary for such a system. Several case studies have been conducted to test and evaluate the resulting integrated system.

The main issue is how to establish a correspondence between how a human perceives spatial and functional aspects of an environment, and what the robot autonomously learns as a map. Most existing approaches to robot map building, or Simultaneous Localization And Mapping (SLAM), use a metric representation of space. Humans, though, have a more qualitative, topological perspective on spatial organization (McNamara, 1986). We adopt an approach in which we build a multi-layered representation of the environment, combining metric maps and topological graphs (as an abstraction over

geometrical information), like (Kuipers, 2000). We extend these representations with conceptual descriptions that capture aspects of spatial and functional organization. The robot obtains these descriptions either through interaction with a human, or through inference combining its own observations (*I see a coffee machine*) with ontological knowledge (*Coffee machines are usually found in kitchens, so this is likely to be a kitchen!*). We store objects in the spatial representations, and so associate the functionality of a location with that of the functions of the objects present there. A core characteristic of our approach is that we analyze each utterance to obtain a representation of the meaning it expresses, and how it (syntactically) conveys that meaning – rather than just doing for example keyword spotting. This way, we can properly handle the variety of ways in which people may express assertions, questions, and commands. Furthermore, having a representation of the meaning of the utterance we can combine it with further inferences over ontologies to obtain a complete conceptual description of the location or object being talked about. This way we can ground situated dialogue in the situational awareness of the robot.

Following (Topp & Christensen, 2005) and (Topp et al., 2006), we talk about *Human-Augmented Mapping* (HAM) to indicate the active role that human-robot interaction plays in the robot's acquisition of qualitative spatial knowledge. In §2 we discuss various observations that independently performed Wizard-of-Oz studies have made on typical interactions for HAM scenarios, and we

indicate which types of interactions we will be able to handle. In §3 we present our approach to multi-layered conceptual spatial mapping and the mechanisms it uses to encode knowledge about spatial and functional aspects of the environment. In §4 we describe the natural language processing facilities that enable the robot to conduct a situated dialogue with its human user about their environment. We present the implementation of our approach in an HRI architecture in §5. In these sections, examples gathered from sample runs at the German Research Center for Artificial Intelligence (DFKI) illustrate the way information is processed in our architecture. §6 presents descriptions of the robot's behavior in test runs carried out at the Royal Institute of Technology (KTH), followed by a discussion of our experiences with the system in §7. The paper closes with conclusions.

## 2. Observations on HAM

Various Wizard-of-Oz studies have investigated the nature of human-robot interaction in HAM. (Topp et al., 2006) discuss a study on how a human presents a familiar indoor environment to a robot, to teach the robot more about the spatial organization of that environment. (Shi & Tenbrink, 2005) study the different types of dialogues found when a subject interacts with a robot wheelchair (while being seated in it). Below we discuss several crucial insights these studies yield.

The experimental setup in (Topp et al., 2006) models a typical guided tour scenario. The human tutor guides the robot around and names places and objects. One result of the experiment is the observation that tutors employ many different strategies to introduce new locations. Besides naming whole rooms ("this is the kitchen" referring to the room itself) or specific locations in rooms ("this is the kitchen" referring to the cooking area), another frequently used strategy was to name specific locations by the objects found there ("this is the coffee machine"). Any combination of these individual strategies could be found during the experiments. Moreover, it has been found that subjects only name those objects and locations that they find interesting or relevant, thus *personalizing* the representation of the environment that the robot constructs.

In the study presented in (Shi & Tenbrink, 2005), the subjects are seated in a robot wheelchair and asked to guide it around using verbal commands. This setup has a major impact on the data collected. The tutors must use verbal commands containing deictic references in order to steer the robot. Since the perspective of the human tutor is identical to that of the robot, deictic references can be mapped one-to-one to the robot's frame of reference. One interesting finding is that people tend to name areas that are only passed by. This can either happen in a 'virtual tour' when giving route directions or in a 'real guided tour' ("here to the right of me is the door to the room with the mailboxes."). A robust conceptual mapping sys-

tem must therefore be able to handle information about areas that have not yet been visited.

Next we discuss how we deal with the above findings, combining information from dialogue and commonsense knowledge about indoor environments.

## 3. Spatial Organization

In order for a robot to be able to understand and communicate about spatial organization, we must close the gap between the different ways humans and robots conceive of spatial entities in their environment.

*Spatial entities*, such as e.g. rooms, areas or floors, are the units of the spatial organization of an (indoor) environment. We assume that spatial and functional aspects of the environment define what the spatial entities are and how one can refer to these entities in situated dialogue. We discuss here our approach to representing the spatial and functional aspects of an environment at multiple levels of abstraction, thus closing this gap. *Spatial aspects* cover the organization of an environment in terms of the geometry, shape and boundaries of connected areas and gateways that together constitute a conception of free and reachable space. *Functional aspects* are higher-order properties that allow or disallow an embodied agent to perform specific actions, such as e.g. passing through a doorway or preparing a meal. In our current approach, we associate functional aspects with an area on the basis of objects present in it. Through dialogue, we can build, query, and clarify these representations, and we point out how they are used in carrying out tasks.




Fig. 1. Snapshot of an online visualization that shows an example of an automatically acquired metric map of a part of the DFKI language technology lab. It shows line features detected in the environment (extended to 3D planes to facilitate viewing) used for SLAM. A navigation graph of interconnected nodes represents free and reachable space. Large red stars indicate doorways and the different colouring of the nodes depicts the topological partitioning of the environment.

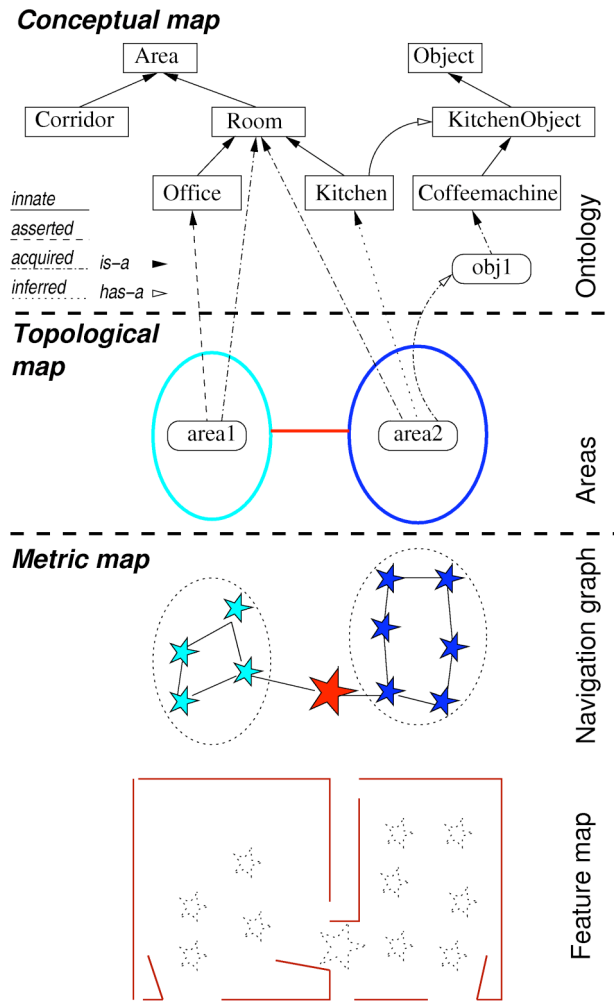


Fig. 2. The three layers of the spatial representation containing a simplified map of an exemplary situation

### 3.1 Representing the environment

The spatial organisation of an (indoor) environment is represented at three levels (Fig. 2).

At the lowest level, we have a *metric map* (Fig. 1), capturing observed spatial structures in the environment with a feature-based representation and establishing a notion of free and reachable space through a navigation graph. The example of Fig. 2 shows line features, which typically correspond to walls. Each map primitive (line features and navigation nodes) is parameterized in world coordinates. A line is for example defined by a start- and an end-point. The metric map is automatically generated from sensor data as the robot moves around the environment. Using features has several advantages. Firstly, they give a compact representation, which, secondly, allows for efficient updates. Among the disadvantages, we find that the map does not explicitly model the free space of the environment as for example an occupancy-grid model would. Only structures that fit the model primitives (e.g. lines) will be captured.

We therefore represent the free space and its connectivity as a *navigation graph*. When the robot moves around, it adds nodes to the graph at the robot's current position if there is not already a node close by in the graph. This

approach is inspired by the notion of 'free space markers', cf. (Newman et al., 2002). Each node is associated with a coordinate in the reference frame of the metric map and thus states that the area around that position was free from obstacles when it was added to the graph. Assuming a mostly static environment, this location is likely to be free also when revisiting it. When the robot travels between nodes, edges are added to the graph to connect the corresponding nodes. We distinguish between two types of nodes: normal nodes and gateway nodes. The gateway nodes (large red stars in Fig. 1 and Fig. 2) encode passages between different areas and typically correspond to doors. In the current implementation, doorways are detected from the laser range data when passing through a narrow opening.

As an intermediate level of abstraction, we have a *topological map*, which divides the navigation graph into areas that are delimited by gateway nodes. There is evidence that humans adopt a topological representation of spatial organization, cf. (Stevens & Coupe, 1978), (Hirtle & Jonides, 1985), and (McNamara, 1986). The topological layer of our map can thus be considered a first approximation of a more humanlike perspective on space.

Fig. 1 shows a real example of a map that the robot has built. The metric map is represented by the lines, which have been extended to pseudo-3D walls to indicate that they typically correspond to walls. The navigation graph is shown as the connected set of stars. The larger (red) stars are gateway nodes and, as can be seen, connect different rooms in this case. The grouping of the nodes in the topological map is illustrated by colouring the nodes. Each area has its own colour.

Finally, we have a *conceptual map* at the top level. In this layer, we store knowledge about names of areas and information about objects present therein. Through fusion of acquired (from sensor data) and asserted (given by the human user) information and innate conceptual knowledge (given in a handcrafted commonsense ontology, cf. §3.2) a reasoner can infer new, additional knowledge. This includes inferences over how a room can be verbally referred to, what kinds of objects to expect in a room, and ultimately a functional understanding of what can be done where and why.

### 3.2 A commonsense ontology of an office environment

Since the robot may have observed only part of an area and the objects therein, and since, as we already pointed out in §2, humans do not necessarily convey complete information about a room, the robot needs to be able to infer knowledge on the basis of only partial information. For this, we use knowledge about spatial and functional properties of an indoor office environment. This conceptual knowledge is modeled in a handcrafted commonsense ontology (Fig. 3). Within this ontology, spatial entities, i.e. instances of the top-level concept *Area*, are categorized on the basis of their spatial and functional properties. Our system makes use of the approach by (Martí-

nez Mozos et al., 2005) to categorize areas into *Corridor* or *Room* based on scans obtained from the laser range sensor in that area, thus considering its spatial aspects. Further distinctions are then made using functional aspects. The subconcepts of *Room* are defined by the instances of *Object* that are found there. This is encoded by the *hasObject* relation. If a given instance of *Room* is related with a specific instance of *Object* by an instance of the *hasObject* relation, this *Room* instance fulfills the conditions for being an instance of the respective specific subconcept of *Room*.

Based on the knowledge representation in the ontology, we use description-logics reasoning to infer general names for rooms, places to look for specific objects, and to resolve linguistically given references to spatial entities (cf. §4.4). Asserted knowledge about locations and objects is derived from structural descriptions of verbal input of a human tutor originating in the communication subsystem (cf. §4.1). Acquired knowledge is derived from laser sensor data via the place classifier or from automatically recognized objects provided by a visual object recognition subsystem (cf. §5.3).

The simplified representation in Fig. 2 shows an example of the interplay between innate, asserted, acquired, and inferred knowledge. The laser-range based classifier assigns *area1* to the concept *Room*. Through interaction with the human user, however, the robot has the more specific information that *area1* can also be referred to as an instance of *Office*. In the next area, *area2*, the automatic classification yields the information that it is of type *Room*. Moreover the camera-based object recognition provides the information that there is an instance of *Coffeemachine* in this area. Using the commonsense knowledge that coffee machines are usually found in kitchens that is encoded in our ontology, it can be inferred that *area2* also instantiates the more specific concept *Kitchen*.

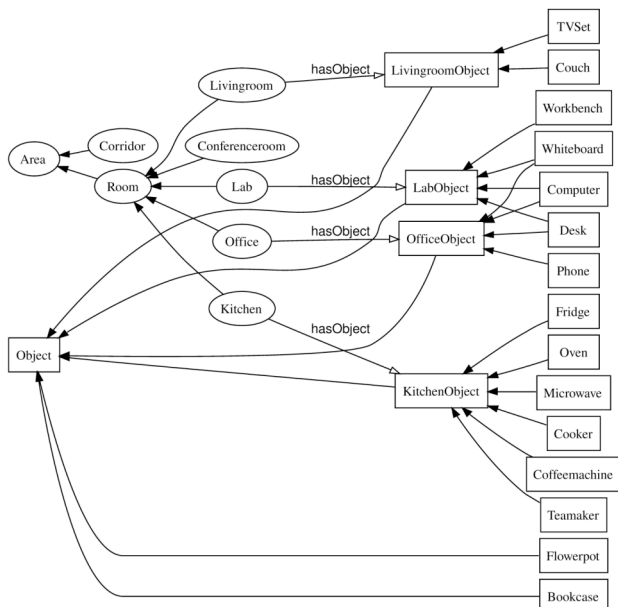


Fig. 3. Commonsense ontology of an office environment. Unlabeled arrows denote the taxonomical is-a relation.

## 4. Situated Dialogue

If robots are to enter the everyday lives of ordinary people, human-robot interaction should minimize the reluctance that people might have towards autonomous machines in their environment. Our natural language communication system accommodates the fact that spoken interaction, *dialogue*, is the most intuitive way for humans to communicate. (Lansdale & Ormerod, 1994) define dialogue as a “joint process of communication,” which “involves sharing of information (data, symbols, context) between two or more parties.” In the context of human-robot interaction (HRI), (Fong et al., 2003) claim that “dialogue, regardless of form, is meaningful only if it is grounded, i.e. when the symbols used by each party describe common concepts.” In the previous section, we have presented our approach to establishing a common conceptual ground for a human-robot shared environment. In this section, we will present the linguistic methods used for natural language dialogue with a robot. We will also address the role of dialogue for supervised map acquisition and task execution.

### 4.1 Deriving the meaning of an utterance

On the basis of a string-based representation that is generated from spoken input through a speech recognition software, a Combinatory Categorical Grammar (CCG) (Steedman & Baldridge, 2003) parser analyzes the utterance syntactically and derives a semantic representation in the form of a Hybrid Logics Dependency Semantics (HLDS) logical form, (Kruijff, 2001) and (Baldridge & Kruijff, 2002). HLDS offers an ontologically richly sorted relational representation of different sorts of semantic meaning: *propositional content* and *intention*. Complex logical forms can be differentiated further by the ontological sort of their intention and their propositional content. Ex. 1–3 show semantic representations of some utterances that would lead to the situation depicted in Fig. 2. Ex. 3 shows the meaning representation for the assertion “This is a bookcase.”

It consists of several, related elementary predicates (EPs). One type of EP represents a discourse referent as a proposition with a handle:  $@_{[B1:thing]}(\mathbf{bookcase})$  means that the referent *B1* is a physical object, namely a bookcase. Another type of EP states dependencies between referents as modal relations, e.g. in Ex. 1 we have  $@_{[I1:region]}(\mathbf{in} \ \& \ \langle Dir:Anchor \rangle(L1: location \ \& \ \mathbf{office}))$ , which means that discourse referent *I1* – an enclosed region – is anchored in a region *L1* being an office. We represent regions explicitly to enable later reference to the region using deictic reference (e.g. “there”). Within each EP we can have semantic features, e.g. the deictic pronoun “this” is characterized as having a visual antecedent that is spatially nearby (**proximal**) the speaker. From these semantic representations, *structural descriptions* of the discourse entities they refer to are constructed. The conceptual map is then updated with the information encoded in those

structural descriptions that can be resolved to spatial entities (Ex. 1) or objects in the environment (Ex. 3).

(1) "We are in the office."  
 @[B1:state](**be**  
 & <Mood>**indicative**  
 & <Restr>(W1: person & **we**)  
 & <Scope>(I1: region & **in**  
 & <Dir: Anchor>(L1: location & **office**  
 & <Delimitation>**unique**  
 & <Number>**singular**)))

(2) "Follow me!"  
 @[F3:action](**follow**  
 & <Mood>**imperative**  
 & <Actor>(R8: hearer & **robot**)  
 & <Patient>(I2: speaker & I))

(3) "This is a bookcase."  
 @[B6:state](**be**  
 & <Mood>**indicative**  
 & <Restr>(T1: thing & **this**  
 & <VisualContext>(O1: visualobject  
 & <Proximity>**proximal**)  
 & <Scope>(C1: thing & **bookcase**  
 & <Delimitation>**existential**  
 & <Number>**singular**)))

A structural description is an HLDS logical form of a nominal phrase – i.e. a syntactic constituent whose head is a noun or a pronoun – that ascribes properties to a discourse referent. The following examples show the structural descriptions that can be derived from the complex logical forms of Ex. 1 and Ex. 3.

(4) @[L1:location](**office**  
 & <Delimitation>**unique** & <Number>**singular**)

(5) @[C1:thing](**bookcase**  
 & <Delimitation>**existential** & <Number>**singular**)

In dialogue analysis, the linguistic meaning of an utterance is related to the current dialogue context, in terms of how it *rhetorically* and *referentially* relates to preceding utterances. The *rhetorical relation* of an utterance indicates how the utterance extends the current discourse – for example, we try to relate an answer to a question that preceded it, to represent what the answer is an answer to. (This plays an important role in handling e.g. clarification questions.) The *referential* relations of an utterance indicate how *contextual references* like definite noun-phrases ("the box") and anaphora ("it") can be related to objects that have been mentioned in preceding utterances. After the utterance is related to the preceding context in this way, an updated model of the dialogue context is obtained in the sense of e.g. (Asher & Lascardes, 2003) and (Bos et al., 2003).

#### 4.2 Mediation of meaning

There are several reasons for why we may want to relate content across different modalities in an HRI architecture. One obvious reason is *symbol grounding*, i.e. the connection of symbolic representations with perceptual or motoric interpretations of a situation, to achieve a situated understanding of higher-level cognitive (symbolic) processes. Achieving such an understanding is an active process. We do not only use the fusion of different content to establish possible connections, but also want it to aid in disambiguating and completing information where and when needed. Finally, relating content may actively trigger processes in a modality (e.g. executing a motor action on the basis of a spoken command) or prime how information is processed (e.g. attentional priming). Altogether this means that we cannot just see content as a symbolic representation without further qualification. We consider content as a tuple that provides a characterization in terms of *intention*, *propositional content*, and a *truth-value*.

An intention reflects *why* the content is provided to other modalities in the architecture. The intention influences what a connection with content in other modalities is expected to yield.

Fig. 4 illustrates an ontology of the types of intentions (top figure) we consider here. The ontology is inspired by theories of *speech acts* ((Searle, 1969) and (Core & Allen, 1997)). We discern commands, assertions, and questions. An assertion attributes a characteristic (*assertion.attributive*). A question can inquire after a fact (*question.factual*), whether a particular state is true (*question.polar*), or whether an agent is capable of performing a certain act (*question.ability*). The intention types in Fig. 4 (top) provide a high-level characterization of the purpose of connecting the content to information in other modalities. When we combine them with the ontological characterization of the propositional content, we get a detailed qualification of what we need to try and relate the content to.

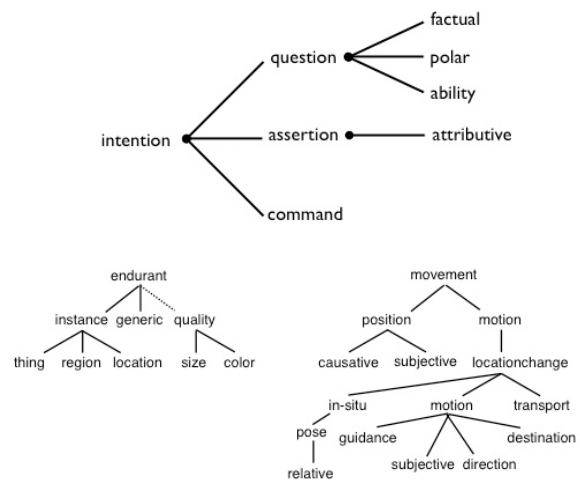


Fig. 4. Two parts of the complex semantic ontology: intention (top) and propositional content (bottom)

Fig. 4 (bottom) gives part of the ontology used for sorting propositional content. It classifies objects (*endurant*) and different types of movement processes (*movement*). Endurants can be physical objects, regions, or locations, and may have qualities such as size or color.

The classification of movement processes is inspired by (Maierborn, 1990) and (Hamp & Feldweg, 1997). Movements are divided into *motion*, and the results of an action as a change of *position*. The latter may concern the agent itself (*subjective*), or regard an object (with the change performed by the agent, *causative*). Motions are movement actions that an agent performs on the spot (*in-situ*), changing pose (*pose.relative*), or that make the agent move to a new position. The latter types of movement can be relative to a person being followed (*guidance*), a subjective frame of reference (*subjective*), a *direction*, or an explicitly given *destination*. Ex. 6 and Ex. 7 illustrate the types.

(6) several typical commands:  
*command. ...*

(a) “go to the laboratory”  
*... movement.motion.locationchange.motion.destination*

(b) “turn to the right”  
*... movement.motion.locationchange.motion.direction*

(c) “follow me”  
*... movement.motion.locationchange.motion.guidance*

(7) “we are in the office”  
*assertion.attributive.endurant.perspective.spatial*

We combine the types of the propositional content with the intentions of Fig. 4. In the examples above, Ex. 6a shows a command to go to a particular destination. Ex. 6b also gives a command. If we change it into “Can you turn to the right?” we get a question after the ability of the agent to turn into a given direction.

Ex. 7 shows the semantic sort of the assertion in Ex. 1.

We create a characterization that includes intention and propositional content so we can determine which modality we need to try and connect this content to. How this connected-to modality should then deal with the provided content is given by the truth-value of the propositional content. The truth-value states how the content can be interpreted against the model of the sensorimotoric or cognitive modality in which the content originates. The interpretation is *dynamic* in that we try to update the model with the propositional content (Muskens et al., 1997). Instead of using a 2-valued truth system, we use a multi-valued system to indicate whether the content was already known in the model (*unknown, known*), and what the result of the update is: *true* if we can update the model, *false* if we cannot, and *ambiguous* if there are multiple ways in which the propositional content can be understood relative to content already present in the model.

To mediate between modalities we represent content using a shared representational formalism, following (Gurevych et al., 2003). We model content as an ontologically richly sorted, relational structure, as described above.

Once we have established the intention, propositional content, and truth value, we can establish *mediation*: We determine to which other modalities we need to establish relations, between the interpretation of the content in the originating modality, and interpretations in those other modalities. Because we determine mediation on the basis of ontological characterizations of content, rather than on its realized form in a modality-specific representation, we speak of *ontology-based content mediation*.

As we already pointed out, mediation can trigger new processes, and result in grounding through information fusion. We keep track of the results of mediation, i.e. the relations between interpretations of content across modalities, by creating *beliefs* that store the handles (identifiers) of the shared representations for the interpretations. We store beliefs at the mediation level. Beliefs thus provide a powerful means for cross-modal information fusion, without requiring individual modalities to commit to more than providing shared representations at the interface to other modalities that enable us to co-index references to interpreted content in individual modalities. For more detail we refer the reader to (Kruijff et al., 2006a).

#### 4.3 Human-Augmented Mapping

In a typical HAM scenario, a human tutor takes the robot on a guided tour of the environment (“follow me!”, cf. Ex. 2 and Ex. 6c). Our robotic system is able to follow its tutor, execute *near navigation* commands (e.g. “turn around!”, “stop!”, cf. (Severinson-Eklundh et al., 2003)), and explore its surroundings autonomously (e.g. “explore the corridor!”, “look around the room!”). These individual behaviors can be freely combined and may be initiated or stopped at any point in time by the human tutor. This mixed control strategy – referred to as *sliding autonomy*, (Heger et al., 2005), or *adjustable autonomy*, (Goodrich et al., 2005) – combines the robot’s autonomous capabilities where appropriate with different levels of tele-control through the human user where needed. However, the human tutor preserves full control over the robot, as he can always stop it or give it new commands.

While thus guiding the robot around, he or she then presents and introduces *locations* (“this is the office.”, cf. Ex. 1 and Ex. 7) and *objects* (“this is the coffee machine.”, cf. Ex. 3). The issue here is how we can use this information to augment the spatial representation.

From language processing, we obtain a representation of the semantics of an utterance (§4.1). Depending on the kind of utterance (e.g. *question, command, assertion*), we decide in what modalities we need to process this content further (§4.2). A prototypical utterance in a HAM scenario makes an *assertion* about the kind of location the

current area is. In this case, we create a structural description (§4.1, Ex. 4 & 5) from the semantics of the utterance, and try to update the conceptual map with the information it contains. The examples below (Ex. 8–10) illustrate a HAM guided tour that would lead to the spatial representation in Fig. 2.

(8) The robot is standing in the office (i.e. area1). It has no initial map of its environment. The human tutor starts a guided tour by asking the robot to accompany him.

H.1 “Come with me!”

R *initialize people following;*  
*direct gaze to the tutor;*  
*acknowledge understanding:*

R.1 “Yes.”

R *start following;*

(9) While the user shows the robot around the room, the robot constructs a metrical map with line features for SLAM and a navigation graph that covers the traveled route. The tutor informs the robot about their location.

H.2 “This is the office.”

R *derive structural description:*  
*@{LI:location}(**office** & <Del>**unique** & <Num>**sg**)*  
*add structural description to conceptual map;*  
*create new instance in ontology:*  
**instance**(area1, Office)  
*acknowledge understanding:*

R.2 “Yes.”

If the human makes an assertion about an object, we anchor the occurrence of the object and its description at the different levels of the spatial representation: in the navigation graph (at the node nearest its position) and in the conceptual map (an instance of the object's type is created and related to the individual that represents the current area), as illustrated in Ex. 10. Note that we do not train the visual object recognizer in a HAM tour. This is done off-line (cf. §6.1).

(10) While still being in the office, the tutor shows the robot the bookcase.

H.3 “This is a bookcase.”

R *derive structural description:*  
*@{B1:thing}(**bookcase** & <Del>**existential** & <Num>**sg**)*  
*add structural description to conceptual map;*  
*create new instance in ontology:*  
**instance**(bookcase1, bookcase)  
**hasObject**(area1, bookcase1)  
*acknowledge understanding:*

R.3 “Yes.”

(11) The tutor then takes the robot to the next room – a kitchen. The robot detects a doorway, creates a gateway node in the navigation graph, and thus creates a new area in the topological map. The place classification classifies the current area (i.e. area2) as *Room*. In

the kitchen, the tutor asks the robot to have a look around. This initiates the automatic vision-based object recognition, which detects a coffee machine. The knowledge about the presence of a *Coffeemachine* in a *Room* is stored in the conceptual map.

H.4 “Have a look around.”

R *initiate autonomous exploration*  
*and turn on visual object recognition;*  
*laser-based place classification classifies area2 as Room:*  
**instance**(area2, Room)  
*vision detects an occurrence of a coffee machine:*

*@{CI:thing}(**coffeemachine**)*

**instance**(obj1, Coffeemachine)

**hasObject**(area2, obj1)

*report the recognition of a coffee machine:*

R.4 “Aha. I see a coffee machine.”

#### 4.4 Answering questions about locations and objects

Given the robot's conceptual map, it is always possible to ask the robot where it thinks it is. If a structural description of the current room has been given before, the robot retrieves this information from the conceptual map (Ex. 11). If the robot has not explicitly been given a general name (such as ‘kitchen’, ‘office’, or ‘lab’) for the current area, the system can try to generate a linguistic expression to refer to the given room. This mechanism makes use of the ontological representation of acquired and innate conceptual knowledge to generate a description (Ex. 12). The description of the area is then returned to the dialogue system, which generates a contextually appropriate utterance to convey the given information (Kruijff, 2006). Depending on whether the referring expression has been retrieved from tutor-asserted information or inferred, either a definite (i.e. <Delimitation>**unique**) or an indefinite (i.e. <Delimitation>**existential**) noun phrase respectively is generated.

(12) The user asks the robot about the locations of the objects it has encountered in the environment. It retrieves the previously asserted information that area1 is “the office” and infers that area2, being a room with a coffee machine, is “a kitchen”.

H.1 “Where is the bookcase?”

R *retrieve structural description for the current area*  
*(area1) from conceptual map:*  
*@{LI:location}(**office** & <Del>**unique** & <Num>**sg**)*  
*generate answer with truth value known\_true:*

R.1 “It is in the office.”

H.2 “Where is the coffee machine?”

R *infer the most specific concept area2 instantiates:*  
**most-specific-instantiators**(area2)  
**returns:** Kitchen

*generate structural description:*

*@{X0:location}(**kitchen** & <Del>**existential** & <Num>**sg**)*

*generate answer with truth value known\_true:*

R.2 “It is in a kitchen.”

If the system fails to find an answer to a question, i.e. the information can neither be retrieved from the conceptual map nor inferred through ontological reasoning, the robot generates a negative answer.

(14)H.3 "Where is the laboratory?"

R *information unavailable;  
generate answer with  
truth value unknown\_false:*

R.3 "I am sorry. I do not know."

#### 4.5 Clarification

Existing dialogue-based approaches to HRI usually implement a master/slave model of dialogue: the human speaks, the robot listens, e.g. (Bos et al., 2003). However, situations naturally arise in which the robot needs to take the initiative, e.g. to clarify an issue with the human. This is one form of mixed-initiative interaction, enabling a robot to recognize when help is needed from a human, and learn from this interaction (Bruemmer & Walton, 2003). A situation that may require is for example when uncertainty arises in automatic area classification: Doors provide important knowledge about spatial organization, but are difficult to recognize robustly and reliably. Clarification dialogues can help to improve the quality of the spatial representation the robot constructs, and to increase the robot's robustness in dealing with uncertain information.

We have extended an approach to processing clarification questions in multi-modal dialogue systems. For space reasons, we refer the reader to (Kruijff et al., 2006b) for technical details. The basic idea is to allow for any modality to raise an issue. An issue is essentially a query for information, which is sent into the architecture. Different modalities, e.g. vision or dialogue, can then respond with a statement that they can handle the query. Once an answer to the query is found, it is then returned to the modality that raised the issue. For example, when mapping is unsure about the presence of a door in a given location, an issue is raised, which is then addressed through interaction with the human. The robot can take the initiative in the dialogue, and phrase a (clarification) question about objects ("What is this thing near me?") or about the truth of a proposition ("Is there a door here?"). Once the dialogue system obtains an answer to the clarification question, both answer and question are provided to the mapping subsystem to resolve the outstanding issue.

#### 4.6 Carrying out tasks

Guiding the robot around an environment is only one step in working with a service robot. The main purpose of a service robot, and of most domestic robots, is to carry out tasks. The multi-level representation of the environment we build up provides an important basis for that. Eventually, we can combine knowledge about what objects are needed to perform particular actions, with the knowledge of where they are. The simplest action to be

performed by a mobile robot is the *go-to task*. The next step in terms of complexity is the *fetch-and-carry task* of locating a specific object or place, going there, possibly fetching the desired object or doing some manipulation with the object in question, and returning.

The current system can be instructed to go to a particular place or object. If the robot knows the location it is sent to, it will just go there. If it has never been shown a place with the respective general name before, it will employ reasoning mechanisms to determine possible locations. If it is sent to an object it has neither been shown nor visually recognized by itself, it will make use of its innate (ontological) knowledge to determine areas that are likely to contain such an object. There, the robot can make use of its autonomous exploration facilities to visually search the area for the desired object.

The current model, however, does not contain functional knowledge about how manipulating and combining objects result in new objects (such as preparing a coffee by placing a cup under a coffee machine and then pressing the start button). As our robot is not equipped with any manipulators (e.g. a gripper or a robotic arm), the physical actions involved in fetching a simple object can only be simulated or replaced by verbally asking for help (Wiltske and Kruijff, 2006).

## 5. Implementation

We have implemented the approach of §3 and §4 in a distributed architecture that integrates different sensorimotor and cognitive modalities. The architecture enables a mobile robot to move about in an indoor environment, and have a situated dialogue with a human about various aspects of the environment.




Fig. 5. The robotic platform



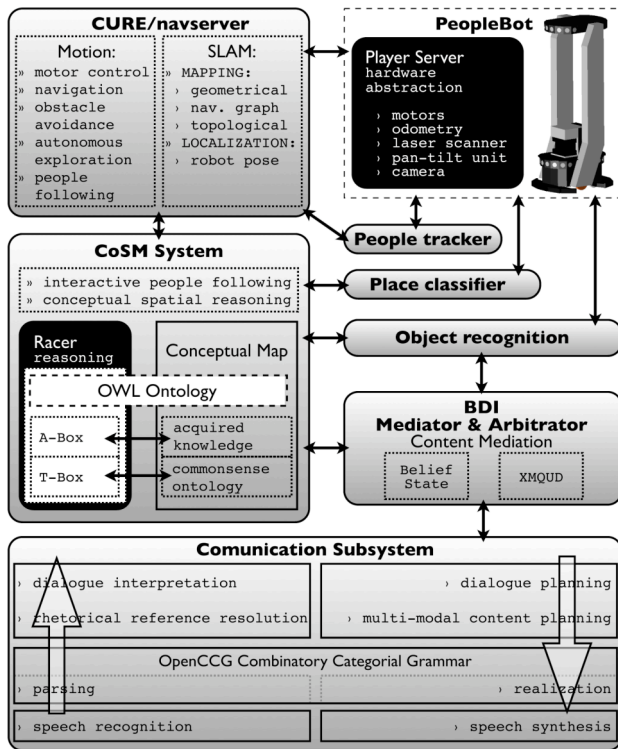


Fig. 6. The architecture

Fig. 5 shows the ActivMedia Pioneer 3 PeopleBot used in the experimental runs at the DFKI language technology lab. It is equipped with a SICK LMS200 series laser scanner, which is the main navigation sensor and is used for building the metric map, performing obstacle detection, tracking people, etc. On top of the robot, there is a Directed Perception pan-tilt unit with a stereo-vision system from Videre Design on it. Bumpers in the front and back are used to detect contact with the environment. The robot hardware is interfaced with using the Player/Stage software. Speech recognition, natural language processing, conceptual spatial reasoning, and people tracking are performed off-board and communicate via wireless Ethernet with the on-board computer.

Fig. 6 shows the relevant aspects of the architecture, with subsystems for situated dialogue, spatial localization and mapping, and visual processing. A BDI-mediator (Belief, Desire, Intention) is used to mediate between subsystems. By this we mean that beliefs provide a common ground between different modalities, rather than being a layer on top of these. Beliefs provide a means for cross-modal information fusion, in its minimal form by co-indexing references to information in individual modalities. The BDI mediator decides what modalities should further process linguistically conveyed information, and how to handle requests for clarifying issues that have arisen. We describe each of these components in more detail below.

### 5.1 The communication subsystem

The communication subsystem consists of several components for the analysis and production of natural utterances in situated dialogue. The purpose of this system is

twofold. Firstly, to take an audio-signal as input, recognize what is being said, and then produce a representation of the contextually appropriate meaning of the utterance. As mentioned before, this then enables combining the conveyed meaning with further inferencing over ontological knowledge. Secondly, to take a representation of meaning *to be conveyed* as input, produce a plan of how the robot can communicate that meaning, and carry out that plan.

The communication subsystem has been implemented as a distributed architecture using the Open Agent Architecture (Cheyer & Martin, 2001). On the analysis side, we use the Nuance speech recognition engine with a domain-specific speech grammar (<http://www.nuance.com>). The string-based output of Nuance is then parsed with an OpenCCG parser. OpenCCG (<http://openccg.sf.net>) uses a combinatory categorial grammar (Baldrige & Kruijff, 2003) to yield a representation of the linguistic meaning for the recognized string/utterance (Baldrige & Kruijff, 2002). These representations are in the same framework used to mediate content between modalities. This enables us to combine linguistically conveyed meaning with further inferences over ontologies.

To produce flexible, contextually appropriate interaction, we use several levels of dialogue planning. Based on a need to communicate, arising from the current dialogue flow or from another modality, the dialogue planner establishes a communicative goal. We then plan the content to express this goal, possibly in a multi-modal way using non-verbal (pose, head moves) and verbal means. During planning, we can inquire the models of the situated context (e.g. dialogue context, visually scene) to ensure the plan is contextually appropriate (Kruijff, 2006). The system realizes verbal content using the OpenCCG realizer. OpenCCG takes logical forms representing meaning as input, and then generates a string for the utterance using the same grammar as we use for parsing utterances (White, 2006). Finally, we synthesize the resulting string using the Mary (<http://mary.dfki.de>) text-to-speech system.

### 5.2 Conceptual Spatial Localization & Mapping

Our method of multi-layered conceptual spatial mapping is handled in two separate modules of our architecture.

#### 5.2.1 The CURE/navserver module for SLAM

The navserver module for SLAM (Simultaneous Localization And Mapping) and robot control is based on the same components that were used in (Jensfelt et al., 2005) and (Folkesson et al., 2005) and is part of the CURE/toolbox software (<http://www.cas.kth.se/CURE>). It creates a metric map that uses lines extracted from the laser-range data as map primitives. The basis for integrating the feature observations is the extended Kalman filter (EKF).

It is also in the navserver module that the navigation graph is constructed and maintained. In the current im-

plementation, the robot adds a new node to the navigation graph when it has moved 1m assuming that there is no old node close by. It builds the topological map automatically from the navigation graph by labeling the nodes with different area identifiers and thus partitions the navigation graph into sets of nodes that correspond to distinct areas in the environment. Our strategy rests on the simple observation that the robot passes a door to move between areas. Whenever the robot passes a door, a node marked as a door is added to the navigation graph and consecutive nodes are given a new area identifier. Currently, door detection is simply based on detecting when the robot passes through a narrow opening. The fact that the robot has to *pass* through an opening removes many false doors that would result from simply looking for narrow openings that appear as valleys in the laser scans. However, this alone will still lead to some false doors in cluttered rooms. A *loop closing* algorithm is used to spot inconsistencies (Kruijff et al., 2006b) arising from falsely recognized doors, and then trigger a clarification dialogue (§4.5).

#### 5.2.2 Ontological Reasoning

The conceptual layer and its links to the lower levels of our map are maintained in the CoSM module. It also provides the link to the communication subsystem by augmenting the topological map with a humanlike conceptual representation of the spatial organization of the robot's environment in an ontology. Ontological reasoning is used to fuse knowledge about types and instances of types in the world. We have built a commonsense ontology of an indoor (office) environment (Fig. 3) as an *OWL ontology*, having *concepts*, *instances* (individuals belonging to concepts) and *relations* (binary relations between individuals). The ontology covers types of locations and typical objects. A priori, as the robot has not yet learnt anything, the ontology does not contain any instances. The robot creates instances as the it discovers its environment (§4.3). For each new area, a new instance of concept *Area* is created. A further distinction between *Rooms* and *Corridors* is provided by the place classification module described in (Martinez Mozos et al., 2005), which is connected to the navserver module. When the robot is in a room, and is shown or visually detects an object, we create a new instance of the corresponding *Object* subconcept, and relate the object's instance and the room's instance using the *hasObject* relation.

We use RACER (<http://www.racer-systems.com>) to reason over *TBoxes* (terminological knowledge / concepts in our ontology) and *ABoxes* (assertional knowledge / instances). We use *assertions* about instances and relations to represent knowledge that the robot learns as it discovers the world. This includes explicit introductions by the tutor or autonomously acquired information. We do not change the TBox at runtime.

If the conceptual map does not contain a structural description that is relevant for the current task (cf. §4.4 and

§4.6), we try to infer the missing information. We use ABox retrieval functions as a first reasoning attempt. The reasoner checks if it can infer that an instance is consistent with the given description. If so, this instance is taken. Else, we use TBox reasoning as a second attempt to resolve uncertainties, e.g. when the robot has not been shown explicitly the occurrence of a relevant object. The robot can thus make use of its a priori knowledge about typical occurrences of objects and use this as a basis for autonomous planning.

Fig. 2 has already briefly sketched how partial information can be fused to infer new concept instantiations.

#### 5.3 Vision

The vision subsystem uses an implementation of SIFT (Scale Invariant Feature Transform) features (Lowe, 2004) for recognizing typical objects in the environment like television sets, coffee machines or bookcases. The system recognizes instances of objects, as opposed to categories, that have previously been shown to the robot and learnt by the visual detection module. The object detection works in an on-line fashion while the robot is moving around. It is turned off, however, as long as the robot is following its human user, who is typically occluding a considerable part of the field of view. When the robot is told to autonomously explore its surroundings, the camera input is used to recognize objects.

#### 5.4 Interactive people following

In order to follow the tutor, we use a laser range based people tracking software (Schulz et al., 2003) that uses a Bayesian filtering algorithm. The people tracker derives robust tracking information of dynamic objects within the robot's perceptual range. Given the tracking data, the people following module calculates appropriate motion commands that are sent to the robot control system to follow the tutor's trajectory, while preserving a socially appropriate distance to the tutor when standing still. The system is *interactive* in that it actively gives the tutor feedback about its state. A pan-tilt-unit with a stereo vision device is moved to always point to the tutor, thus conveying the robot's user awareness. Note that the cameras are not used to track the user, but serve only the purpose of providing *gaze feedback*.

#### 5.5 Situational and functional awareness

We currently investigate how the information encoded in the multi-layered conceptual spatial representation can be used for a smarter, human- and situation-aware behavior. As one aspect of this, the robot should exploit its knowledge about objects in the environment to move in a way that allows for successful interaction with these objects.

For instance, when following a person, the robot should make use of its knowledge about doors in the environment, such that it recognizes when the person wants to perform an action with the door. As actions that are performed in a doorway or with the door itself potentially

require a wide space, e.g. for swinging or sliding open the door, for letting people pass, or for stepping past the door opening to grab the door handle, it is crucial that the robot adjusts its actions accordingly. A failure to understand such a situation could, for example, lead the robot to a position where it traps the user in the doorway that he or she was trying to close.

In the current implementation, we opt for the robot to increase the distance it keeps to the user when it detects that the user approaches a door and to decrease it again when it detects that the user left the area. In this way, as the robot does not stop tracking and following the person, the people following behavior stays smooth and intuitive for the user.

## 6. Case studies

We have carried out several experimental runs of the complete integrated system in the DFKI language technology lab and at the 7<sup>th</sup> floor of the CAS building at KTH. The system used for the experiments at KTH differs from the system used at DFKI (cf. §5) in that it features a Canon VC-C4 pan-tilt-zoom camera instead of a stereo camera. The length of the studies ranges between several minutes and more than half an hour. The robot was operated by one tutor at a time using verbal commands only. No telecontrol or Wizard-of-Oz techniques were used. In the runs, the robot was guided through several situations visiting several rooms. The tutors, lab employees familiar with the system, were equipped with a Bluetooth headset connected to the automatic speech recognition software. The software modules of our integrated system were running in real-time on several laptops that were interconnected via a wireless network. The onboard computer of the robot, which is also equipped with a wireless adapter, was running the hardware abstraction drivers. The laptop running the speech recognizer was placed on the bottom deck of the robot platform to ensure a reliable Bluetooth connection to the headset.



Fig. 7. The tutor activating the robot



Fig. 8. The tutor interacting with the robot. The trashbin close to the table, which created the illusionary doorway can be seen.

There were no specific tasks defined for the experimental runs. They were rather used to test and evaluate the overall functionality of the integrated system. The processing of typical human-robot dialogues of the sample runs at DFKI has been illustrated by examples in the previous sections. The following paragraphs illustrate the behavior of the robot during the experimental runs at KTH. Videos illustrating sample runs with our system are available at <http://www.dfki.de/cosy/www/media>.

### 6.1 Training phase

We had collected the training data for the laser-range bases place classifier beforehand and trained the classifier off-line. The SIFT-based object recognition had also been trained off-line on different objects that were later used in the experimental runs. These objects included among others a TV set, a couch, and a bookcase. The acquisition of the training data was not part of the experimental runs, but the results were part of the innate knowledge that the robot had in the beginning of the runs.

### 6.2 Activating the robot

In the beginning of the test runs (Fig. 7) the robot was standing in the corridor close to a wall-mounted charging station. The speech recognition software was running, but operating in a quiet mode. In this mode, the robot does not react to any verbal command except for the explicit command to “wake up” and listen to the tutor. This way, the dialogue system is not forced to process interaction between the tutor and the other experimenters, which would lead to many falsely recognized utterances and an erratic behavior of the robot. By uttering “partner, wake up” the tutor activates the full speech processing facilities of the robot, which is thus ready to follow the tutor.

### 6.3 Following the tutor

Next, the tutor commanded the robot to follow him. While following the tutor from the corridor to a room, the

navserver reliably detected the doorway when it was passed through. The place classifier reliably classified the corridor correctly. Due to the tutor's presence at a close distance to the laser sensor, the scans obtained in the room were sometimes misclassified as corridor. Especially when there were long laser beams that reached beyond open doors to the side of the robot and very short beams reflected by the tutor's legs in front of the robot, the resulting scan profiles resembled the ones typically encountered in corridors.

#### 6.4 Clarification dialogues

By placing a trash bin at a distance of about 80 cm from a table and guiding the robot through this narrow opening, we created an illusionary doorway in the middle of the room (cf. Fig. 8). After passing through it, the robot assigned the subsequent new navigation nodes a new area identifier. The robot then reached a previously visited node with a different area identifier. This inconsistency reliably triggered the clarification dialogue as presented in (Kruijff et al., 2006b). The robot asked the user "is there a door here?". Using the online viewing software (cf. Fig. 1), the tutor was able to understand the robot's question. Upon the user's denial of the question, the map of the robot was corrected by relabeling the false doorway and merging the two areas.

#### 6.5 Inferring new concepts

So far, the visual object recognition had been turned off. When the user commanded the robot to "have a look around," the visual recognition was turned on and subsequently detected the couch and the television set (Fig. 9). Combining this acquired knowledge from vision with the acquired knowledge that the current area is of type *Room*, the inference mechanisms of the conceptual spatial mapping subsystem yielded the information that the current area can be categorized as a *Livingroom*. This was verified in the experiments by the user asking the robot where it is twice, once before turning on vision and once after. In the former case, the robot replied "I am in a room," reflecting that it only had acquired type information and no tutor-asserted knowledge about how to refer to the current area. In the second case, the robot correctly answered "I am in a living room," which showed that the conceptual reasoning worked as intended.

## 7. Experience

Our main experience with the implemented system is that there are a couple of principal behaviors needed for HAM. If we want a human to guide a robot around an environment, then the robot must be able to (a) follow the human, (b) use information it gets from the human to augment its map, (c) take the initiative to ask the human for clarification; and (d) we need to be able to verify, and correct, what the robot has (not) understood. Where is the system successful, and where is it not?



Fig. 9. "Aha. I see a television."

a) Although people tracking and following works fairly smoothly, the robot tends to lose track when the human e.g. passes around a corner. We are now studying how to predict the *path* where a tracked human is going, to overcome this problem and to reduce misclassifications of static objects as dynamic (due to laser data noise). We have also found that having a notion of what human behavior to expect is important: when a human moves to open a door, the robot should not follow the human behind the door, but go through it. The robot needs to reason over functionality of regions and objects in the environment to raise such expectations. We are currently investigating how we can make use of the knowledge that the robot has about its environment to allow for a *smarter* behavior in situations like mentioned before.

b) The question here is not just whether the robot can use information from the human – there is also the issue of how easy or difficult it is for the human to convey that information to the robot in the first place. In our grammar, we have *lexical families* that specify different types of syntactic structures and the meaning they convey, and *lexical entries* specifying how words belong to specific lexical families. This way we can specify many ways in which one can convey the same information (*synonymy*). Dialogue can thus be more flexible, as there is less need for the human to know and give the precise formulation (controlled language).

c) Clarification often concerns aspects of the environment which need to be explicitly referred to, e.g. "Is there a door *here*?" The difficulty lies in generating deictic references with a robot with a limited morphology. Although we can generate spatial referring expressions, non-verbal means would be preferable. However, body- and head-pose may not be distinctive enough. We may thus have to drive to a place (the "HERE") to make the deictic reference explicit, while avoiding disturbing the interaction.

d) Because we have reliable speech recognition (recognition rate is >90%), misunderstanding is primarily a semantic issue. This raises two main questions. First, how does the human understand that the robot understood

what was said, without asking the robot? Various systems have the robot repeat what it has just heard. We have not done this; the robot only indicates whether it has understood (“yes”/“okay”/“no”). We have not experienced problems with this, but we are investigating now more explicit non-verbal cues for grounding feedback (e.g. gaze). Second, we need to study what types of misunderstanding may occur in HRI for HAM, and to what extent they may have a *relevant* effect on the robot's behavior. This is an issue we now investigate.

e) The fact that the reasoner in the current implementation of the system works in a strictly monotonic way makes it impossible to clarify overgeneralizations of the robot's inferences. If, for instance, an office worker keeps a teamaker for personal use in his or her office and the robot detects this with its object recognition software, it will infer that this office can also be referred to as “kitchen”. We currently investigate how such overgeneralizations can be ruled out in clarification dialogues and how the reasoning mechanisms have to be adapted to prevent negative statements like “no, this is not a kitchen” from making the A-Box knowledge inconsistent.

f) When the user presents the robot with new objects, e.g. “this is the coffee machine,” the robot should follow the gaze of the person or look for pointing gestures in order to be able to acquire a visual model of the objects referred to. A related interesting question here is how the robot can make sure that it in fact is the right object that it has found without using a monitor to interact with its tutor.

## 8. Conclusions

We presented an HRI architecture for human-augmented mapping and situated dialogue with a human partner about the environment they share. We discussed the multi-level representations we build of the environment, including spatial organization and functional aspects (based on salient objects present in areas). The system uses autonomous mapping, visual processing, human-robot interaction, and ontological reasoning to construct structural descriptions with which the multi-level representations are annotated. The approach has been fully implemented, and helps bridging the gap between robot and human conceptions of space. We showed its functionality, inspired by independently performed Wizard-of-Oz studies, on several running examples. For future research we want to study more detailed spatial organizations of regions and objects within rooms, to create 3-dimensional representations.

## 9. Acknowledgements

The research reported in this article was supported by the EU FP6 IST Cognitive Systems Integrated Project *Cognitive Systems for Cognitive Assistants* “CoSy” FP6-004250-IP.

## 10. References

- Asher, N. & Lascarides, A. (2003). *Logics of Conversation*, Cambridge University Press, New York, NY, USA
- Baldrige, J. & Kruijff, G.-J. M. (2002). Coupling CCG and Hybrid Logics Dependency Semantics, Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL 2002), pp. 319–326, Philadelphia, PA, USA, 2002
- Baldrige, J. & Kruijff, G.-J. M. (2003). Multi-modal combinatory categorial grammar, Proceedings of the 10th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2003), Budapest, Hungary, 2003
- Bos, J., Klein, E. & Oka, T. (2003). Meaningful conversation with a mobile robot, Proceedings of the Research Note Sessions for the 10th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2003), pp. 71–74, Budapest, Hungary, 2003
- Bruemmer, D. J. & Walton, M. (2003). Collaborative tools for mixed teams of humans and robots, Proceedings of the Workshop on Multi-Robot Systems, Washington, D.C., USA, 2003
- Cheyner, A. & Martin, D. (2001). The Open Agent Architecture. *Journal of Autonomous Agents and Multi-Agent Systems*, Vol. 4, No. 1, pp. 143–148
- Core, M. & Allen, J. (1997). Coding dialogs with the DAMSL annotation scheme, Proceedings of the AAAI Fall Symposium on Communicative Action in Humans and Machines, Boston, MA, USA, 1997
- Folkesson, J., Jensfelt, P. & Christensen, H. I. (2005). Vision SLAM in the measurement subspace, Proceedings of the 2005 IEEE International Conference on Robotics and Automation (ICRA 2005), Barcelona, Spain
- Fong, T., Nourbakhsh, I. & Dautenhahn, K. (2003). A survey of socially interactive robots. *Robotics and Autonomous Systems*, Vol. 42, pp. 143–166
- Fritz, M., Leibe, B., Caputo, B. & Schiele, B. (2005). Integrating representative and discriminant models for object category detection, Proceedings of the International Conference on Computer Vision (ICCV05), Beijing, China, 2005
- Goodrich, M. A., Olsen, D. R., Crandall, J. W. & Palmer, T. J. (2001). Experiments in adjustable autonomy, Proceedings of the IJCAI-01 Workshop on Autonomy, Delegation and Control: Interacting with Intelligent Agents, Seattle, WA, USA, 2001
- Gurevych, I., Porzel, R., Slinko, E., Pfleger, N., Alexandersson, A. & Merten, S. (2003). Less is more: Using a single knowledge representation in dialogue systems, Proceedings of the HLT-NAACL Workshop on Text Meaning, Edmonton, Alberta, Canada, 2003
- Hamp, B. & Feldweg, H. (1997). GermaNet – A lexical-semantic net for German, Proceedings of the ACL Workshop on Automatic Information Extraction and

- Building of Lexical Semantic Resources for NLP Applications, Madrid, Spain, 1997
- Heger, F. W., Hiatt, L. M., Sellner, B., Simmons, R. & Singh, S. (2005). Results in sliding autonomy for multi-robot spatial assembly, Proceedings of the 8th International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS 2005), Munich, Germany, 2005
- Hirtle, S. C. & Jonides, J. (1985). Evidence for hierarchies in cognitive maps. *Memory and Cognition*, Vol. 13, pp. 208–217
- Jensfelt, P., Ekvall, S., Kragic, D. & Aarno, D. (2005). Integrating SLAM and object detection for service robot tasks, IROS 2005 Workshop on Mobile Manipulation: Basic Techniques, New Trends and Applications, Edmonton, Alberta, Canada, August 2005
- Kelleher, J. D. & Kruijff, G.-J. M. (2005). A context-dependent model of proximity in physically situated environments, Proceedings of the 2nd ACL-SIGSEM Workshop on the Linguistic Dimensions of Prepositions, Colchester, Essex, UK, 2005
- Kruijff, G.-J. M. (2001). *A Categorical-Modal Logical Architecture of Informativity: Dependency Grammar Logic & Information Structure*. PhD thesis, Faculty of Mathematics and Physics, Charles University, Prague, Czech Republic
- Kruijff, G.-J. M. (2006). Dependency grammar. *The Encyclopedia of Language and Linguistics*, 2nd edition, Elsevier Publishers
- Kruijff, G.-J. M., Kelleher, J. D. & Hawes, N. (2006a). Information fusion for visual reference resolution in dynamic situated dialogue, In: *Perception and Interactive Technologies (PIT 2006)*, Andre, E., Dybkjaer, L., Minker, W., Neumann, H. & Weber, M. (Eds.), Springer Verlag, 2006
- Kruijff, G.-J. M., Zender, H., Jensfelt, P. & Christensen, H. I. (2006b). Clarification dialogues in human-augmented mapping, Proceedings of the 1st ACM Conference on Human-Robot Interaction (HRI 2006), Salt Lake City, UT, USA, March 2006
- Kuipers, B. (2000). The Spatial Semantic Hierarchy. *Artificial Intelligence*, Vol. 119, pp. 191–233
- Lansdale, M. & Ormerod, T. (1994). *Understanding Interfaces*. Academic Press, London, UK
- Lowe, D. G. (2004). Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, Vol. 60, No. 2, November 2004, pp. 91–110
- Maienborn, C. (1990). Position und Bewegung: Zur Semantik lokaler Verben, Technical Report, IWBS-Report No. 138, IBM Stuttgart, Stuttgart, Germany
- Martínez Mozos, O., Stachniss, C. & Burgard, W. (2005). Supervised learning of places from range data using adaboost, Proceedings of the 2005 IEEE International Conference on Robotics and Automation (ICRA 2005), pp. 1742–1747, Barcelona, Spain
- McNamara, T. (1968). Mental representations of spatial relations. *Cognitive Psychology*, Vol. 18, pp. 87–121
- Muskens, R., van Benthem, J. & Visser, A. (1997). Dynamics, In: *Handbook of Logic and Language*, van Benthem, J. & ter Meulen, A. (Eds.), Elsevier, 1997
- Newman, P., Leonard, J., Tardós, J. & Neira, J. (2002). Explore and return: Experimental validation of real-time concurrent mapping and localization, Proceedings of the 2002 IEEE International Conference on Robotics and Automation (ICRA 2002), pp. 1802–1809, Washington, D.C., USA, 2002
- Schulz, D., Burgard, W., Fox, D. & Cremers, A. B. (2003). People tracking with a mobile robot using sample-based joint probabilistic data association filters. *International Journal of Robotics Research*, Vol. 22, No. 2, pp. 99–116,
- Searle, J. R. (1969). *Speech Acts: An Essay in the Philosophy of Language*, Cambridge University Press, Cambridge, UK
- Severinson-Eklundh, K., Green, A. & Hüttenrauch, H. (2003). Social and collaborative aspects of interaction with a service robot. *Robotics and Autonomous Systems*, Vol. 42, pp. 223–234
- Shi, H. & Tenbrink, T. (2005). Telling Rolland where to go: HRI dialogues on route navigation, Proceedings of the Workshop on Spatial Language and Dialogue (5th Workshop on Language and Space), Delmenhorst, Germany, 2005
- Steedman, M. & Baldridge, J. (2003). Combinatory Categorical Grammar (draft 4.0)
- Stevens, A. & Coupe, P. (1978). Distortions in judged spatial relations. *Cognitive Psychology*, Vol. 10, pp. 422–437
- Topp, E. A. & Christensen, H. I. (2005). Tracking for following and passing persons, Proceedings of the 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2005), pp. 70–76, Edmonton, Alberta, Canada, August 2005
- Topp, E. A., Hüttenrauch, H., Christensen, H. I. & Severinson-Eklundh, K. (2006). Acquiring a shared environment representation, Proceedings of the 1st ACM Conference on Human-Robot Interaction (HRI 2006), pp. 361–362, Salt Lake City, UT, USA, March 2006
- White, M. (2006). Efficient realizations of coordinate structures in combinatory categorical grammar. *Research on Language and Computation*, Vol. 4, No.1, 2006, pp. 39–75
- Wilske, S. & Kruijff, G.-J. M. (2006). Service robots dealing with indirect speech acts, Proceedings of the 2006 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2006), Beijing, China, October 2006