

A Dialog-Driven Process of Generating Route Directions

Kai-Florian Richter^{a,b,*} Martin Tomko^{a,c} Stephan Winter^a

^a*Department of Geomatics, University of Melbourne, Australia*

^b*Transregional Collaborative Research Center SFB/TR 8 Spatial Cognition,
Universität Bremen, Germany*

^c*Department of Geography, University of Zürich, Switzerland*

Abstract

Humans adapt the instructions provided in route directions to the assumed spatial knowledge of the receivers; the majority of route directions is provided to wayfinders with at least partial spatial knowledge of the environment. However, today's navigation systems assume no a-priori knowledge. Most of current research addresses this by exploring means to personalize assistance through the capture of knowledge about individual users. Accordingly, such systems require an extended learning phase. In this paper, an approach to adaptive route directions based on a combination of turn-by-turn directions and destination descriptions is presented. This approach does not rely on information on a wayfinder's previous knowledge. Instead, a wayfinder can adjust the type and detail of the presented information via dialog. The paper focuses on the problem formalization and its algorithmic realization; the approach is generic with respect to the modality of the actual dialog (e.g., verbal or key-press based computer interfaces). The paper provides a contribution towards non-static, adaptive route direction services.

Key words: route directions, destination description, turn-by-turn directions

* Kai-Florian Richter, FB03, Universität Bremen, PO Box 330 440, 28334 Bremen, Germany; phone: +49 421 218 64283; fax: +49 421 218 64239

Email addresses: richter@sfbtr8.uni-bremen.de (Kai-Florian Richter), martin.tomko@geo.uzh.ch (Martin Tomko), winter@unimelb.edu.au (Stephan Winter).

1 Introduction

Wayfinders always have some knowledge of their environment, either from previous experience, or acquired indirectly through communication (e.g., in the form of external representations, such as maps), or inferred from previous experiences of similar environments. Even without any help of route directions, wayfinders can apply wayfinding strategies to find their destination. Route directions are useful for wayfinders only to the extent they support the decisions in their wayfinding strategies, i.e., to the extent they shorten the travel to the destination, lower the cognitive workload of wayfinding, or improve the wayfinders' decisions with respect to other cost functions.

If the wayfinder asks another person, a local *expert*, for route directions, the communication situation gives a rich resource of context, which is used by the expert to identify information assumed to be relevant for the wayfinder. The communication situation gives, for example, cues about the prior knowledge of the wayfinder. Today's wayfinding services, such as web-based route planners and car navigation systems, ignore prior knowledge of the wayfinders—and hence, may come up with irrelevant or even patronizing detail. To avoid this, in current research the capture of individual wayfinders' profiles is used to infer the extent of a-priori spatial knowledge and to adapt the content of the resulting route directions (Cheng et al., 2004; Patel et al., 2006). This is a process of personalization that is time-consuming and still far from practical applications as we will argue below.

The aim of this paper is to identify a *generic* way of automatically generating adaptive route directions; solving the problem of generating flexible route directions that make use of prior knowledge of individual wayfinders. We achieve this by combining turn-by-turn directions (Richter, 2007a) and destination descriptions (Tomko, 2007). As we show, an adaptive generation of route directions free from specific assumptions about prior knowledge of the individual wayfinder can only be established via dialog. Hence, the paper focuses on how to conclude the meaning of wayfinders' requests and how they influence the generation of a response. By this way, adaptive route directions emerge from the in-situ interaction between system and wayfinder.

Several research questions have to be answered to specify and implement such a process. First, if no assumptions on prior knowledge shall be made, how can a route direction generation algorithm decide on appropriate types of responses, and appropriate references in the responses? And if route directions are constructed piecewise, on request, how can it be guaranteed that the pieces have proper transitions to each other, and give a complete and consistent picture of the route? Finally, how do these dialog-driven route directions relate to existing cognitively motivated models of route directions, such as the destination

descriptions of Tomko (2007) or the aggregations in turn-by-turn directions of Richter (2007a)?

The hypothesis of this paper is that two accessible categories of context, namely the *route context* and the *dialog context*, are sufficient to generate adaptive route directions, i.e. route directions that are arbitrarily coarse or detailed, depending on the wayfinders' current needs. The route context defines the wayfinder's location along a route. Specifically, three different locations can be distinguished: the origin, the destination, and a location en route. The route context is derived from information received from a positioning device. The dialog context keeps track of previous requests made by the wayfinder, the references provided by the system so far, and the kind of instructions the system currently delivers. Both contexts in their combination allow inferring the intention of a wayfinder's request.

In particular, this paper studies the commonalities and differences in destination descriptions and turn-by-turn directions, and their transitions in route directions containing both forms of communication. It will be shown that when seeking guidance, wayfinders ask one of two questions: either a '*where*' question or a '*how*' question. It will also be shown that there is a formal dependency between the type of the question and the context considered (the combined route context and dialog context). Since we know how to respond to a 'where' question—by destination descriptions (Tomko, 2007)—and to a 'how' question—by aggregated turn-by-turn directions (Richter, 2007a)—the remaining task is to identify the transitions between both forms of responses. The paper proposes an algorithm that solves this problem, and demonstrates the behavior of the algorithm in a range of wayfinding scenarios.

The paper makes a significant contribution towards the development of non-static, dialog-driven route direction services, which has relevance in human-computer interaction as well as cognitive engineering. The proposed algorithm considers the environment and the context of a request to tailor route directions for the specific information needs of an individual wayfinder. It does so without capturing any information about the personal history of the wayfinder, such as prior spatial knowledge, this means it is truly generic. Implementing this algorithm as part of next-generation wayfinding services will help to improve their usability as well as market acceptance.

The paper is structured as follows. In the next section, we introduce the problem in detail and review relevant background information for the presented research. Section 3 introduces the underlying models for turn-by-turn directions and destination descriptions. In Section 4, it is shown that dialog facilities are needed to combine both approaches, which results in a model for dialog-driven route directions (Section 4.1); its algorithmic realization is detailed in Section 4.2. Finally, the model is applied to some example scenarios in Section

5, and in Section 6 conclusions and future work are presented.

2 Problem Description and Related Work

2.1 *The Wayfinding Problem*

Wayfinding is a purposive, directed, motivated activity to follow a route from origin to destination (Golledge, 1999). According to Montello (2005), it reflects the cognitive processes going on during navigation—as opposed to *locomotion*, which covers the activities of the sensory and motor system.

When wayfinding in an environment, for example a city, people acquire knowledge about that environment; they form an “image of the city” (Lynch, 1960). This image is acquired over time. With increasing exposure, people remember previously visited places and landmarks, learn routes between them, and integrate the knowledge into a more coherent survey representation (Siegel and White, 1975; Montello, 1998). The layout of an environment, i.e. its structure, influences how easily and integrated its mental representation is formed by people (Dogu and Erkip, 2000; Werner and Long, 2003). The structure of an environment, and people’s familiarity with an environment, also influences the strategies they employ to find their way around an environment (Conroy Dalton, 2001; Hochmair and Raubal, 2002; Hölscher et al., 2006).

2.2 *Forms of Route Directions*

People seek external support when they are not certain about the way to take. This support, either provided by another human or an assistance service, provides additional information for wayfinding in the form of route directions. Route directions are task-oriented specifications of the actions to be carried out to reach a destination (Denis, 1997; Tversky and Lee, 1999). Generally, they may be provided in-advance (a priori) before route following starts; in this case, usually information about the complete route is given. Instructions may also be given incrementally while already following a route (in situ); then usually only information about the next turning action due is provided (Habel, 2003; Richter, 2007b).

Human route directions reflect the direction giver’s knowledge about an environment: upon being asked for directions, humans activate the spatial knowledge of the route to be described, identify the relevant information, structure this information, and communicate it to the requester (Denis, 1997; Lovelace

et al., 1999; Allen, 2000). Since a route is a temporal and spatial sequence of relevant features and actions, route directions describing these features and actions are also presented sequentially in order of their occurrence. Usually, references in route directions are given from the perspective of the wayfinder; they focus on those points along the route where the wayfinder has to decide on the further way to take—the *decision points* (Daniel and Denis, 1998). To anchor these actions in space, humans frequently refer to landmarks (Denis, 1997; Tversky and Lee, 1999).

Human route directions show two different structures: answers to ‘where’ questions (Shanon, 1983), or answers to ‘how’ questions (Lovelace et al., 1999). Which structure (or which combination of both structures) humans choose as an answer to an assistance request depends on the context (Porzel et al., 2006). Generally, a ‘where’ question asks for the location of the destination; a ‘how’ question asks for procedural information concerning the way to get to that destination. An answer to the former is termed *destination description* (Tomko, 2007); it provides a description of the destination’s location relative to other features of the environment. A typical example for a destination description is describing the location of your home to a taxi driver at the airport (“Neukirchstraße, please. That’s at Torfhafen, opposite to the exhibition center”). An answer to a ‘how’ question is termed *turn-by-turn direction*; it provides an instruction for every turn from the origin to the destination. A typical example is describing the way to your home to a friend from another city arriving at the train station (“Leave the station through the southern exit; cross the fairground, keeping the exhibition center to your right; turn right behind it; walk straight on; cross the street at the traffic lights, and then cross the bridge to your left.”).

As these examples show, destination descriptions are useful for wayfinders who have at least partial knowledge of the environment, whereas turn-by-turn directions are needed by those wayfinders who do not know any specifics about the environment. Or, put differently, destination descriptions suffice for known parts of an environment, whereas turn-by-turn directions are needed to cover unknown parts. Since for most wayfinders there are known and unknown parts in any environment of geographical scale, wayfinding assistance systems should be able to combine both kinds of instructions. Imagine, for example, a first-time visitor to a new environment: she will quickly build up route knowledge as partial knowledge of the environment. Or imagine a local: he knows the environment from long-term use, but may be challenged by dynamic changes in the environment, such as road closures (construction sites, congestion) or a reversion of one-way directions.

However, an assistance system requires information about the wayfinder’s spatial knowledge in order to decide when a switch between the forms of directions is required and appropriate. That is, an assistance system needs to decide

whether a wayfinder is familiar with the current (part of the) environment and, accordingly, destination descriptions suffice, or whether they are in an unknown (part of the) environment and the wayfinder requires turn-by-turn directions. Such decisions can only sensibly be made if the system has information about the individual’s spatial knowledge (Patel et al., 2006; Schmid and Richter, 2006; Srivinas and Hirtle, 2007).

Cheng et al. (2004) present an approach to a car navigation system that distinguishes between known, potentially known, and unknown parts of an environment, derived from a user’s previous navigation behavior. For known parts of an environment, only intermediate destinations, such as highway exits, are given; for unknown environments detailed turn-by-turn directions are provided. The system may initiate a dialog with the user if it cannot determine whether the current part of the environment is known or unknown; however, users cannot initiate a dialog themselves. User initiated dialog is part of the system presented by Hurtig and Jokinen (2006); their PDA-based navigation system for public transport allows a user requesting more detailed information using simple spoken commands (such as ‘navigate more’). However, their focus is on the integration of multi-modal input and on disambiguating references, using wayfinding mainly as application scenario.

While systems that adapt to users’ previous knowledge would deliver unobtrusive and truly adaptive directions, they need to be personalized to each individual user. Thus, such a service cannot be used on an ad-hoc basis, as it needs to acquire and integrate data on each user’s knowledge. This has to be done over the course of multiple wayfinding interactions, while reasoning about which places and parts of the environment have become familiar to a wayfinder. Consequently, this personalization is a time-consuming and usually imperfect process.

3 Generation of Route Directions

The underlying representation for generating route directions is a graph of the street network. A route is represented as a directed path in that graph, connecting the node representing the origin with the node representing the destination. In today’s commercial systems, the change of heading at a decision point, i.e. the angle between incoming and outgoing edge at a node, is typically described using qualitative direction relations, such as ‘turn left’ or ‘turn right’. These instructions are annotated with the name of the street a wayfinder has to turn into; the way to the next turning event is specified by metric distances (e.g., ‘247m’, ‘10km’).

While generating route directions this way is computationally simple and ef-

ficient, it is significantly different from the way humans produce route directions; the resulting directions are cognitively demanding to process for a human wayfinder. Accordingly, several approaches to generating more human-like, cognitively motivated route directions exist (e.g., Maaß, 1994; Tversky and Lee, 1999; Habel, 2003; Dale et al., 2005; Klippel et al., 2005). These approaches take into account elements that are not part of the network itself, for example, the elements identified by Lynch (1960). Especially the identification (e.g., Raubal and Winter, 2002; Elias, 2003) and integration (e.g., Caduff and Timpf, 2005; Hansen et al., 2006) of landmarks play a crucial role. In the presented approach to dialog-driven route directions, the generation of turn-by-turn directions is based on the approach to *context-specific route directions* (Richter and Klippel, 2005; Richter, 2007a), which is further detailed in Section 3.1.

Underlying the generation of destination descriptions is also a route, i.e. a directed path in a graph. However, these descriptions do not provide information on each segment of the route (i.e., each edge), but instead specify the location of the destination (Tomko and Winter, 2006b; Look and Shrobe, 2007). Destination descriptions refer to features of the environment in the surroundings of the destination that are prominent; prominence is derived from hierarchies of features. For example, different landmarks may be ordered according to their (visual, structural, and cognitive) salience, ranging from the ATM at the corner to well-known landmarks, such as the Eiffel tower (cf. Sorrows and Hirtle, 1999). The generation of destination descriptions is further explained in Section 3.2.

3.1 *Generating Turn-By-Turn Directions*

In the model developed in Section 4, turn-by-turn directions are generated using the approach to context-specific route directions (Richter, 2007a). Context-specific route directions account for environmental characteristics and a route’s properties; they adapt to the current action to be taken in the current surrounding environment, i.e. respect for the route context. They are termed context-specific because of the explicit adaptation to the structure and function in wayfinding (Klippel, 2003). A computational process, called GUARD, has been developed for generating context-specific route directions (Richter, 2007a). GUARD stands for Generation of Unambiguous, Adapted Route Directions. Route directions generated by GUARD unambiguously describe one route to the destination, with instructions adapted to environmental characteristics. Figure 1 provides an overview of the generation process.

GUARD works on a graph representing an environment’s street-network. This graph is annotated with information on landmarks, for example, their location

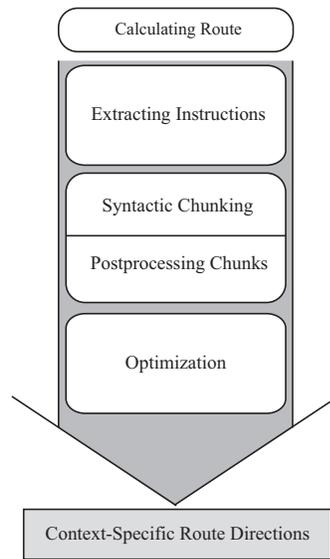


Fig. 1. Overview of GUARD, the generation process for context-specific route directions.

and shape. The generation of context-specific route directions is a four-step process. In the first step, for every decision point of the route, all instructions that unambiguously describe the route segment to be taken are generated, resulting in a set of possible instructions for each decision point. To this end, the generation process makes use a systematics of route direction elements. Specifically, GUARD employs references to different types of landmarks in generating instructions—namely point, linear, and areal landmarks—whose role in the route directions depend on their location relative to the route (Hansen et al., 2006; Richter, 2007c).

Next, GUARD performs *spatial chunking* (Klippel et al., 2003). Chunking reflects an important mechanism of the way humans produce route directions: the combination of several instructions for consecutive decision points into a single instruction, for example, “turn right at the third intersection” or “turn left at the supermarket.” This chunking covers steps two and three of the generation process. GUARD is flexible with respect to the principles used in these steps. For example, it allows integrating the chunking principles presented by Klippel et al. (2003) or Dale et al. (2005).

Finally, in the fourth step of GUARD, the actual context-specific route directions are generated. Here, from all possible instructions those that best describe the route are selected. As this is realized as an *optimization* process, ‘best’ depends on the chosen optimization criterion. Just as with the chunking principles, GUARD is flexible with respect to this criterion, i.e. it is possible to implement and combine different optimization criteria. Optimization results in a sequence of chunks that cover the complete route from origin to destination. However, due to the aggregation of instructions performed in chunking,

instructions for some decision points may only be represented implicitly; thus, reducing the communicated information. Examples of context-specific route directions can be found in Section 5, where the generation of dialog-driven route directions is demonstrated. The complexity of generating these directions is dominated by the optimization step; it depends on the number of decision points covered by the route.

3.2 Generating Destination Descriptions

A destination description, as defined by Tomko (2007), is a referring expression uniquely describing a destination of a route in a given urban environment. Such descriptions are provided during inferential communication to a recipient with at least partial a-priori knowledge of the environment. Destination descriptions consist of a hierarchically ordered set of references to prominent features. The references are communicated from most prominent to least prominent, thus specifying the destination's location with increasing level of detail. The references only locate the destination, they do not provide any instructions on the route to get there.

Relevance of an utterance, as defined by Sperber and Wilson (1986), is a function of the cognitive effect of the utterance on the hearer, and of the cognitive effort required to process the utterance. For destination descriptions, the relevance of a reference to a spatial feature is proportional to the feature's prominence (derived from the experiential hierarchy), and inversely proportional to its distance from the start and destination of the route. The more prominent an element of the environment is, the less effort is required to relate the reference to one's own mental representation of the feature referred to. Similarly, the distance from the referent increases the cognitive effort, as the ambiguity of the interpretation of the reference increases. While the Eiffel Tower, for example, is the most prominent landmark in Paris, it is not relevant when traveling in the outskirts of the city due to its huge distance to the route. On the other hand, an ATM, even though not generally prominent, may be highly relevant when located at one of the route's decision points.

References are retrieved from a hierarchically organized, integrated dataset of heterogeneous elements of the city (districts, streets and landmarks), organized in order of decreasing prominence. This dataset can be based on the same graph as used for context-specific route directions (see last section). Additionally, regions representing districts and the reference regions of landmarks need to be represented. The measures of prominence selected for the various types of elements of the city are cognitively motivated, and facilitate the estimate of the wayfinders' shared experience of prominence. Administrative hierarchies have previously been tested on districts (Tomko and Winter, 2006a), while

novel approaches are used to derive such hierarchies for landmarks (Winter et al., to appear 2007) and streets (Tomko et al., to appear 2007).

A set of rules guides the reference selection process, selecting the most relevant references from the candidate set. Each selected reference contributes to the dialog context. The selection process recursively adapts to the changing context after the selection of every reference. As the interpretation of each consecutive reference will occur in the context of the previous one, the selection must adapt in a similar way. Additional parameters, such as preference for references to features that are part of, or located along the route are also integrated in the selection rule set.

In the model of destination descriptions used, a physical and linguistic co-presence is assumed. Upon request for a description of a specific destination, the references constituting the resulting referring expression are retrieved. First, the system calculates a route from the start to the destination, based on an arbitrary choice of a cost function, for example, the shortest path. Second, a candidate set of features covering the destination or found in its proximity is created. From there, the references for the destination description are retrieved, evaluating the relevance of features in a candidate set. The result of the selection process is an ordered set of references, containing only a small number of references that fits in the working memory of the recipient. In special cases when the destination is a highly prominent feature unique in the given route context, a direct reference may be sufficient (e.g., a reference to the destination Eiffel Tower is sufficient to reach it for any taxi driver anywhere in Paris). The complexity of calculating destination descriptions is largely determined by selecting references from the ordered sets of features. Thus, the complexity is proportional to the depth of the hierarchical structure storing the dataset. It does not depend on the length of the route or the number of intersections passed along it.

4 A Generic Model of Dialog-Driven Route Direction Generation

There are significant differences between destination descriptions and turn-by-turn directions. Accordingly, there is no straightforward combination of both approaches. While both refer to the street network and prominent features—the landmarks—of an environment, the underlying elements used for calculating instructions differ. Generation of destination descriptions is based on named streets as constitutive elements (e.g., “the hotel is in Turnbull Alley, off Spring Street”). Turn-by-turn directions build on decision points (intersections) as their basic elements (e.g., “turn left before the Old Treasury building (at the corner Collins Street / Spring Street); turn left into Turnbull Alley”). Even more, neither does an increasing refinement of destination descriptions

turn them into turn-by-turn directions, nor does an increasing abstraction of turn-by-turn directions turn them into destination descriptions. Even on the finest level of granularity, destination descriptions only locate the destination and do not provide instructions on intermediate decisions along the route. And while it is possible to employ approaches to turn-by-turn directions in such a way that not for every single route segment an instructions needs to be generated (e.g., Höök, 1991; Richter, 2007b), these directions still contain descriptions of intermediate locations along a route. Furthermore, for a sensible combination of both types of route directions, an assistance system needs to know in which parts of an environment to use which form of route directions, i.e. it requires information about a wayfinder’s spatial knowledge (see Section 2).

4.1 Dialog-Driven Route Directions

To get around the problem of needing to know what the wayfinder knows, dialog facilities are introduced to the wayfinding assistance service. The wayfinder is by default presented with destination descriptions, assuming that the environment is known, and can request further, more detailed information by using the dialog facility if the currently presented information does not suffice. This way, an assistance service is achieved that is usable in an ad-hoc manner, i.e. without the need to train the service first. Using dialog to negotiate known and unknown parts of an environment is a straightforward, unobtrusive solution to the problem at hand.

In the following, we present a generic model of such a dialog-driven generation process of route directions; the concrete modality of interaction (e.g., natural language processing or simple button presses) is not important for this model. Route directions are generated in situ. Thus, the dialog facility must be designed such that a wayfinder can request further information anywhere along a route. The presented solution can be implemented to run self-contained on a mobile device, i.e. the device does not need on-line access. It needs a positioning mechanism; however, it is sufficient to correctly position wayfinders on the street segment they are currently on, i.e., the part between two intersections. To properly handle requests from the wayfinder, the system needs to keep track of two different kinds of context: the route context and the dialog context. Depending on the location along the route, requesting more detailed instructions may have different consequences, as will be further explained below.

Figure 2 shows an overview of the dialog-driven generation process. Upon a request, the system generates destination descriptions for the requested destination, answering a ‘where’ question. The wayfinder may now start traveling

towards that destination or may request more detailed information, this way posing a ‘how’ question. Such a request results in turn-by-turn directions being calculated to the next prominent location along the route and presented to the wayfinder. As soon as the wayfinder reaches this location, the original destination descriptions are checked and possibly re-calculated and again presented as wayfinding assistance. This is a recurring process: further requests result in more turn-by-turn directions. If no further prominent location is found along the route ahead—as it happens for example with requests made from the last elements of the destination description—turn-by-turn directions lead to the actual destination.

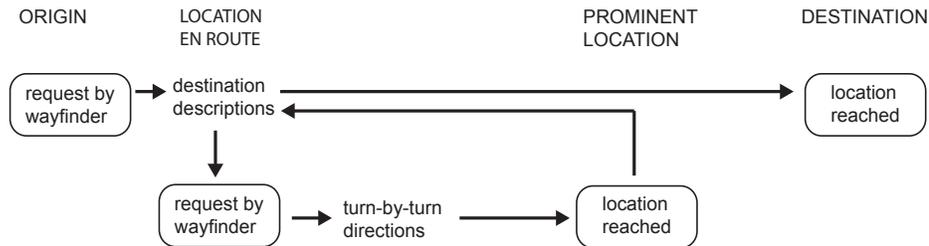


Fig. 2. Overview of the dialog-driven generation of route directions. The system generates destination descriptions as default instructions, but a wayfinder may request more detailed instructions anywhere along the route.

Prominence of locations is defined by the hierarchies of spatial features as explained in Sections 2 and 3.2. Generating route directions this way may result in different patterns of instructions, displayed in Figure 3. The two basic cases—pure destination descriptions and pure turn-by-turn directions—are shown on the top and bottom, respectively. Pure turn-by-turn directions emerge if a wayfinder requests further information at each prominent location along a route. Such turn-by-turn directions guide a wayfinder from intermediate destination to intermediate destination; the system does not generate in-advance directions for the complete route (as is the case in the approach by Richter, 2007a).

The underlying reason for guiding a wayfinder to the next prominent location is that people familiar with an environment only need assistance locally to navigate unknown or complex parts; reaching a prominent location gets them back on track. That is, the irrelevant and patronizing amount of instructions given to expert wayfinders is minimized for the costs of a novice wayfinder having to request assistance more than once. In this approach, prominence is not based on an individual’s spatial knowledge, but defined by a threshold level of the hierarchies of environmental features. These hierarchies capture the (assumed) shared knowledge of people familiar with an environment; prominence arises from environmental properties, not from individual experiences. In the following subsection, an algorithm for the dialog-driven generation of route directions is presented.

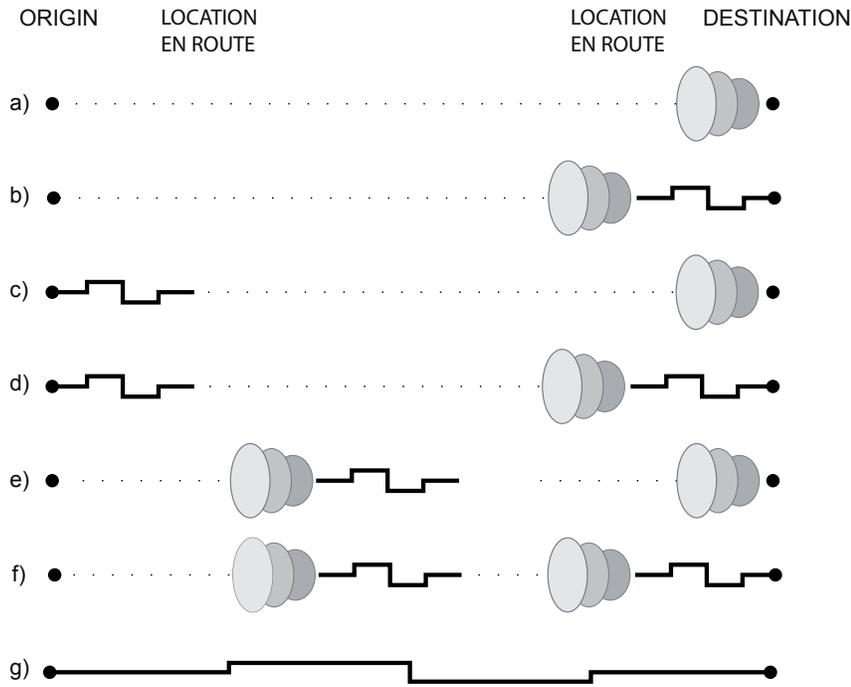


Fig. 3. Different patterns of instructions that may emerge from a dialog-driven generation. The displayed patterns only show canonical cases; arbitrarily different ones may arise. Patterns a) and g) show the two basic cases: pure destination description and pure turn-by-turn direction.

4.2 An Algorithm for the Dialog-Driven Generation of Route Directions

Dialog-driven route directions start with a destination description for the wayfinders' destination, which is presented to them when wayfinding starts. If there are no more requests by the wayfinder during the trip, this destination description is all the information that is provided. The current wayfinder's location is tracked throughout the complete trip, though. If the wayfinder requests additional information there is a switch to *turn-by-turn mode* ('tbt' in Algorithm 1) and turn-by-turn directions from the wayfinder's current location to the next intermediate destination are calculated. This intermediate destination is the actual destination if the region corresponding to the element on the finest level of granularity in the destination description has already been reached. Otherwise, the next prominent location gets calculated as intermediate destination. The turn-by-turn directions are presented to the wayfinder and it is tracked whether they reached the intermediate destination. As soon as they do, there is a switch back to *destination description mode* ('dd' in Algorithm 1) and the destination description from the newly reached location to the final destination is recalculated and presented to the wayfinder again (destination descriptions are not recalculated if a wayfinder already reached the region of finest granularity). This process, which is summarized in Algorithm 1, continues until the wayfinder reaches the final destination. It ensures that

the wayfinder can request further information any time, but also ensures that it only produces information if requested, thus avoiding the communication of irrelevant information to the wayfinder. Tracking a wayfinder’s position and waiting for requests can be assumed to have constant complexity. Thus, the complexity of the algorithm depends on the complexity of generating destination descriptions and turn-by-turn directions, respectively.

Algorithm 1 The dialog-driven route direction generation process.

```

mode ← tbt
while not destination reached do
  if mode = tbt and not destination_region reached then
    mode ← dd
    calculate destination description from current location
    destination_region ← region of finest granularity
  end if
  while not (user request or destination reached) do
    update user location
    listen for user request
  end while
  if user request then
    mode ← tbt
    if destination_region reached then
      intermediate_destination ← destination
    else
      intermediate_destination ← next prominent location
    end if
    calculate turn-by-turn direction to intermediate_destination
    while not intermediate_destination reached do
      update user location
    end while
  end if
end while

```

Previewing the scenario used in the next section, being at the airport and requesting an address in inner-city Melbourne may result in a destination description such as “Turnbull Alley is in the CBD, off Spring Street, opposite the Parliament.” Along the way to Melbourne’s inner city, the wayfinder’s location is tracked. A request by the wayfinder results in turn-by-turn directions being calculated from their current location to the next prominent location along the route. To determine this location, a path between current location and destination has to be calculated, for example the shortest path. Additionally the level of prominence of the street the wayfinder is currently on has to be stored. This prominence value is the threshold value that determines the next prominent location. As argued below, this next destination must at least match the prominence of the current location en route, and it has to be along the route. If such a location can be found, it is set as an intermediate

destination to which turn-by-turn directions are generated. If no such location exists on the remaining route to the destination, the wayfinder is guided by turn-by-turn directions all the way to their destination. For example, upon leaving the highway onto Bell Street, the wayfinder may request further assistance. The shortest path from current location to destination is calculated and Sydney Road is identified as a street with a prominence level higher than Bell Street's. Accordingly, turn-by-turn directions to this street are generated: "turn right before the railway tracks, turn left at the end of the street, then turn right at the traffic lights (onto Sydney Road)."

Since turn-by-turn directions use decision points as their basic elements, streets are used as the environmental features that determine a prominent location, namely the intersections of two or more streets. For every decision point, the level of prominence of each street passing through it is checked. If one of them has a prominence value of at least the threshold value, this decision point is set to be the intermediate destination of the turn-by-turn directions. The prominence of the street on which turn-by-turn directions are requested is used as threshold value because it can be assumed that this level of prominence reflects the wayfinder's knowledge about the street-network for the whole neighborhood. This heuristic reflects the theory of the skeleton in the cognitive map of Kuipers et al. (2003): the structure of the street network influences which streets people travel on frequently, which increases these streets' prominence in a positive feedback loop, resulting in a skeleton of well-known streets covering an environment. Betweenness centrality, which is the fundamental measure used to determine the experiential prominence of streets (Tomko et al., to appear 2007), is a structural measure of the street network's graph that reflects a street's likelihood of being traveled through. Thus, streets on the same level of granularity can be assumed to be on the same level of experience of wayfinders according to the skeleton theory. If, for some reason, this heuristic fails and a wayfinder does not recognize a street chosen as next prominent location turn-by-turn directions can be requested again, thus, keeping the wayfinder on track.

The other elements used in generating destination descriptions—landmarks and districts—cannot be used as prominent locations. Districts are not suitable as they do not allow for identifying single known points along a route. In principle, landmarks could be used. However, just because a landmark is prominent, the streets in its neighbourhood do not necessarily have to be prominent as well (cf. Siegel and White, 1975). Therefore, we also refrain from using landmarks. Algorithm 2 summarizes the determination of the next prominent location.

Algorithm 2 Determining the next prominent location.

```
prominence_threshold  $\leftarrow$  prominence of street corresponding to current location
route  $\leftarrow$  shortest path between current location and destination
intermediate_destination  $\leftarrow$  NIL
current_decision_point  $\leftarrow$  first decision point of route
while not (intermediate_destination or destination reached) do
  prominence_level  $\leftarrow$  max(prominence level of intersecting streets at current_decision_point)
  if prominence_level  $\geq$  prominence_threshold then
    intermediate_destination  $\leftarrow$  current_decision_point
  end if
  current_decision_point  $\leftarrow$  next decision point of route
end while
if not intermediate_destination then
  intermediate_destination  $\leftarrow$  destination
end if
```

5 Application of the Model

In this section, the application of the model in generating route directions is demonstrated. Throughout the examples, wayfinders with different spatial knowledge are assumed; this illustrates how the model tracks the route and dialog context and adapts the kind and detail of spatial information communicated in the generated route directions according to a wayfinder's current needs. The following scenario is used: arriving at Melbourne's international airport (Tullamarine), the respective wayfinders want to make their way to Turnbull Alley, which is in Melbourne's Central Business District (CBD) (see Figure 4).

First, the generation of destination descriptions for this wayfinding task is illustrated, followed by the generation of turn-by-turn directions. Then, several examples of the combination of both kinds of instructions will be given. Note that the model for dialog-driven route directions does not produce actual verbal output (as neither of the underlying approaches produce such output), therefore, the verbal instructions presented in the following should be seen as a possible externalization of the calculated directions that make the examples more readable. References to building names are included for readability reasons as well; in an application producing turn-by-turn directions for wayfinders not knowing the area, these would need to be replaced by a description of the building.

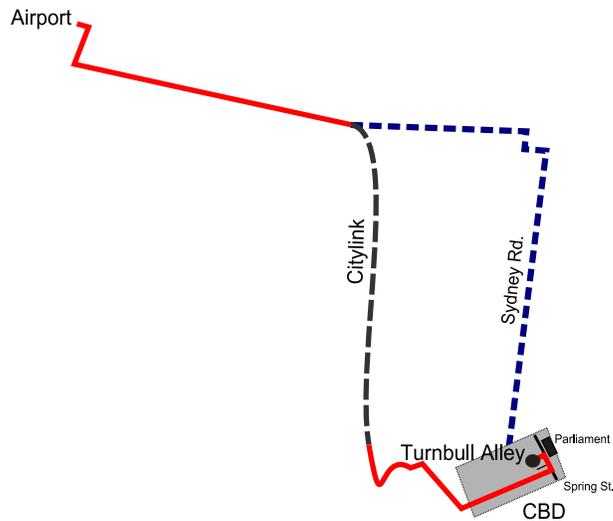


Fig. 4. Sketch-map showing the example trip from Melbourne’s Tullamarine Airport to Turnbull Alley in the CBD.

Generation of destination descriptions In this first scenario, the wayfinder is a local, i.e., a person with knowledge of the structure of the city. In the model, initially route context is set to ‘origin’ and the dialog context is ‘where.’ Accordingly, upon stating the destination location (Turnbull Alley), a destination description is computed as described in Section 3.2. The shortest path between origin and destination is calculated, which is used to guide the search for suitable referring expressions in the hierarchies of addressable features. A relevance threshold is set that determines up to which level features are considered to be relevant.

In this example, the district Turnbull Alley is located in (the CBD) is identified as being the most relevant. The pre-conditions that must hold in order to create a reference to this district are checked next (cf. Tomko and Winter, 2006a); especially whether it is identical or neighbored to the origin’s district, in which case no reference could be generated due to their small topological distance. In our example, however, “CBD” is a valid reference. The selection then continues recursively from within the region “CBD”, more precisely from the point where the calculated route enters this region. Thus, the reference “CBD” and the newly determined origin set the dialog context for selecting the next reference. Turnbull Alley is further contained in the region corresponding to the landmark “Parliament” which also has a relevance above the threshold and satisfies the condition of topological distance. Therefore, “Parliament” is added as the next reference. Finally, the only other feature above the relevance threshold is the street connecting the last reference (“Parliament”) to the destination (“Turnbull Alley”), which is “Spring Street.”

Hence, the destination description contains the following references, ordered from coarse to fine granularity: CBD, Spring Street, Parliament. This set

of references uniquely identifies the location of the destination and provides prominent spatial features in the proximity as references. This description is communicated to the wayfinder, who does not need further assistance, corresponding to the pattern in Figure 3a. The verbalization of this set of references may look as follows:

“(Turnbull Alley) is in the CBD, off Spring Street, opposite the Parliament.”
This type of route description might be helpful for a taxi driver.

Generation of turn-by-turn directions Imagine a wayfinder that is a tourist who has never been to Melbourne before and rented a car at the airport. This wayfinder requires turn-by-turn directions throughout the entire trip. However, upon entering the destination, again by default the destination description is generated. Since this is not helpful, a request for further information is issued. That is, the route context is ‘origin’, the dialog context, however, switches to ‘how’, and turn-by-turn directions are generated.

Again, as a first step the shortest path between origin and destination is determined; the prominence of the street the wayfinder is currently on is stored as threshold for determining the next prominent location. Being in the car park of the airport, this threshold is low. Next, starting from the current location for each decision point in the route it is checked whether any of the intersecting streets has a prominence value above the current prominence threshold. In this case, this holds true for the highway leading from the airport to the city. Thus, using GUARD turn-by-turn directions are generated with this freeway as destination.

Upon reaching the freeway, the route context is ‘location en route’ and the dialog context returns to ‘where.’ Since the wayfinder is unfamiliar with all of Melbourne, again a request for further information is issued, setting the dialog context to ‘how.’ And again, after determining the next prominent location, turn-by-turn directions are generated from the current location to that location. This pattern repeats itself until the wayfinder reaches the destination. In total, the wayfinder has to request further information three times; the following instructions are provided during the trip (each block represents the instructions generated in each request):

From the car park, follow the exit signs.

Turn right at the car park exit.

Follow the street until you enter the freeway (pass the McDonald’s and the ABC childcare).

Follow the freeway until Essendon Airport (after passing under a freeway).

Continue on the freeway and exit at Docklands Highway (passing the red Melbourne Gateway sculpture).

Turn left at the next two intersections.

Turn right.

Follow the railway tracks until you pass Southern Cross Station.

Turn left after Southern Cross Station (into Collins Street).

Turn left before the Old Treasury building.

Turn left into the second street at Parliament building (into Turnbull Alley).”

Unfamiliar destination Next, examples that combine destination descriptions and turn-by-turn directions are considered. In the first of these examples, the wayfinder has only coarse knowledge of the area where the destination is. Upon requesting directions, destination descriptions are determined as explained above. These suffice to get the wayfinder to the CBD and to locate prominent Spring Street. When turning into Spring Street, however, the wayfinder requests further assistance. This requires the system to switch into turn-by-turn mode: even though near the destination, route context is still ‘location en route’, and dialog context switches to ‘how.’

Turn-by-turn directions from the wayfinder’s current location to the next prominent location are calculated. As the wayfinder is already near the destination, the destination region is already reached and the final destination is set as next prominent location (according to Algorithm 2). Getting into Turnbull Alley requires a left turn from Spring Street; Parliament building is identified as a landmark suitable to indicate this turn (cf. Richter, 2007c). Since there are two streets leading to the left at Parliament building, Turnbull Alley is labeled as the second street for disambiguation. At all intersections in-between, the wayfinder needs to continue straight. Accordingly GUARD combines them with the required left turn into a single instruction in the chunking step. The wayfinder receives the following turn-by-turn directions to reach the destination:

“Turn left into the second street at Parliament building.”

Unfamiliar origin This example covers another canonical case of dialog-driven route directions (Figure 3c). While the wayfinder is familiar with the destination location, the area around the origin, the airport, is unknown. This may, for example, happen to a resident of Turnbull Alley that has never navigated in the airport area. Accordingly, after receiving the destination description, the wayfinder immediately requests further information. As in Example 2, this results in a switch to turn-by-turn mode and the next prominent location is calculated; while being still at the origin. The first street that has a higher prominence value than the airport’s car park—where the trip starts—is the freeway running into the city. Thus, turn-by-turn directions to reach this freeway are generated and communicated to the wayfinder:

“From the car park, follow the exit signs.

Turn right at the car park exit.

Follow the street until you enter the freeway (pass the McDonald's and the ABC childcare)."

When the wayfinder enters the freeway, dialog context switches to ‘where’ and destination descriptions are re-calculated. As the wayfinder is still in the region of the airport, the same situation as before holds, i.e. region of origin and destination are not adjacent. Hence, the same set of references as before is identified and communicated to the wayfinder.

Unfamiliar location en route As in the initial example on generating destination descriptions, the wayfinder is familiar with the way to take and at first does not need further assistance. However, on the trip, the wayfinder decides to exit the freeway before entering the Citylink (a tollway) at the last free exit. This takes them to an unknown part of the city and, accordingly, the wayfinder requests more detailed instructions. Route context is ‘location en route’, dialog context switches to ‘how’. A new shortest path to the destination is calculated and the next prominent location along that path is determined. The wayfinder is currently on Bell Street, which is a fairly prominent street in Melbourne. However, Sydney Road, a street of similar prominence level is only a few blocks away; this road is set as next prominent location. The resulting instructions are:

"Exit to the left to Bell Street.

Turn right before the railway tracks.

Turn left at the end of the street.

Turn right at the traffic lights (onto Sydney Road)."

Upon entering Sydney Road, the dialog context switches to ‘where’ and destination descriptions are re-calculated. But, as in the example above, region of origin and destination are not adjacent, i.e. the destination description remains the same.

6 Conclusions and Future Work

Dialog-driven route directions are a generic way of automatically generating adaptive route directions. They bridge two different forms of spatial communication—turn-by-turn directions and destination descriptions. Both forms represent common means to communicate route knowledge to wayfinders. While easily combined in natural communication, until now these two forms have not been integrated in one computational model. In this paper, we developed such a model. To this end, we illustrated its basic constituents—the generation of destination descriptions and of turn-by-turn directions, and presented

an algorithm that covers the generation process of dialog-driven route directions. The presented solution is generic with respect to the used modality. We demonstrated the application of the model in several wayfinding scenarios.

Previous approaches to adaptive route directions rely on personalization of presented information through the acquisition and storage of user profiles. In contrast, the model to dialog-driven route-directions allows wayfinders to adapt the level of spatial information provided to their own spatial knowledge, without the necessity to implement personalization and customization measures, such as storing the users' tracking history beyond the duration of the trip. This reduces issues of privacy that may make users object to using such a service

The examples presented in Section 5 illustrate how route directions produced by the presented model communicate the required information in a concise way that does not patronize a wayfinder with irrelevant information. The destination descriptions initially presented are brief, focusing on the identification of the destination's location. Although this initial information may not be sufficient for a wayfinder, the references provided in the destination description are reproduced in the turn-by-turn directions (with the exception of references to districts as, for example, "*CBD*"). Accordingly, the switch from destination descriptions to turn-by-turn directions does not require the wayfinder to process irrelevant references. A small number of instructions in the turn-by-turn directions is sufficient to guide the wayfinder to the next prominent location. Thus, overall, the model for dialog-driven route directions needs to communicate only little spatial information to get a wayfinder to the destination (compared to internet route-planners, for example); and a good part of it is only produced on demand.

The interaction of the user with the system is minimal; distracting the user from other tasks, such as driving, can, thus, be reduced. In our example, the extreme case of pure turn-by-turn directions required only three interaction steps (requests). In an actual system, there could be an option for overriding it and requesting turn-by-turn directions for the complete route in the initial interaction step, this way further reducing the number of interactions. Further experiments will elicit the amount of additional information that usually needs to be provided in turn-by-turn directions. This will allow to adapt the dialog-based model to various situations, such a route directions for drivers, emergency services or pedestrians.

The model adapts well to situations where wayfinders need to be guided back after a wrong turn (either by asking them to go back or by re-calculating the route). However, it is important to note that getting the wayfinder back on track is only sensibly available in turn-by-turn mode where the wayfinder is supposed to follow a specific route and, thus, their progress towards the

destination can be unambiguously monitored. In destination-description mode no specific route is communicated and, accordingly, it is difficult to judge whether the wayfinder is making progress towards the destination or is lost. As long as the wayfinder does not indicate to being lost—by requesting more information—the system respects the autonomy of the wayfinder, which is an aspect we see as a contribution to user adaptation.

In this paper, the focus is on the problem formalization and the algorithmic aspects of dialog-driven route directions. It is an approach that is generic with respect to the modality of the dialog, i.e. to the actual interface used for interacting with a system producing dialog-driven route directions. Accordingly, issues of human-computer interfaces or issues of a concrete mode of communication (e.g., speech recognition or natural language processing) are not addressed in this paper (for these aspects, see, e.g., Wahlster et al., 2001; Krüger et al., 2004). Different interfaces are possible, from a simple button causing the system to switch into turn-by-turn mode to a true verbal dialog. Currently, work on an example interface using natural language has been taken on. Such an example interface will allow evaluating users' satisfaction and performance with dialog-driven route directions.

Acknowledgments

Part of this work is supported by the Transregional Collaborative Research Center SFB/TR 8 Spatial Cognition which is funded by the Deutsche Forschungsgemeinschaft (DFG).

References

- Allen, G. L., 2000. Principles and practices for communicating route knowledge. *Applied Cognitive Psychology* 14 (4), 333–359.
- Caduff, D., Timpf, S., 2005. The landmark spider: Representing landmark knowledge for wayfinding tasks. In: Barkowsky, T., Freksa, C., Hegarty, M., Lowe, R. (Eds.), *Reasoning with mental and external diagrams: computational modeling and spatial assistance - Papers from the 2005 AAAI Spring Symposium*. Menlo Park, CA, pp. 30–35.
- Cheng, H., Cavedon, L., Robert, D., 2004. Generating navigation information based on the driver's route knowledge. In: *Proceedings of the Coling 2004 Workshop on Robust and Adaptive Information Processing for Mobile Speech Interfaces*. Geneva, Switzerland, pp. 31–38.
- Conroy Dalton, R., 2001. The secret is to follow your nose — route path selection and angularity. In: Peponis, J., Wineman, J., Bafna, S. (Eds.),

- Proceedings of the Third International Space Syntax Symposium. University of Michigan, Ann Arbor, pp. 47.1–47.14.
- Dale, R., Geldof, S., Prost, J.-P., 2005. Using natural language generation in automatic route description. *Journal of Research and Practice in Information Technology* 37 (1), 89–105.
- Daniel, M.-P., Denis, M., 1998. Spatial descriptions as navigational aids: A cognitive analysis of route directions. *Kognitionswissenschaft* 7, 45–52.
- Denis, M., 1997. The description of routes: A cognitive approach to the production of spatial discourse. *Cahiers Psychologie Cognitive* 16 (4), 409–458.
- Dogu, U., Erkip, F., 2000. Spatial factors affecting wayfinding and orientation — a case study in a shopping mall. *Environment and Behavior* 32 (6), 731–755.
- Elias, B., 2003. Extracting landmarks with data mining methods. In: Kuhn, W., Worboys, M., Timpf, S. (Eds.), *Spatial Information Theory. Foundations of Geographic Information Science. International Conference COSIT*, Springer, Berlin, pp. 375–389, LNCS 2825.
- Golledge, R. G., 1999. *Wayfinding Behavior — Cognitive Mapping and Other Spatial Processes*. John Hopkins University Press, Baltimore, USA.
- Habel, C., 2003. Incremental generation of multimodal route instructions. In: *Natural Language Generation in Spoken and Written Dialogue. AAAI Spring Symposium 2003*, Palo Alto, CA.
- Hansen, S., Richter, K.-F., Klippel, A., 2006. Landmarks in OpenLS - a data structure for cognitive ergonomic route directions. In: Raubal, M., Miller, H., Frank, A. U., Goodchild, M. F. (Eds.), *Geographic Information Science - Fourth International Conference, GIScience 2006*. Springer, Berlin, pp. 128–144, LNCS 4197.
- Hochmair, H., Raubal, M., 2002. Topologic and metric decision criteria for wayfinding in the real world and the WWW. In: *Spatial Data Handling. SDH'02*, Ottawa, Canada.
- Hölscher, C., Meilinger, T., Vrachliotis, G., Brösamle, M., Knauff, M., 2006. Up the down staircase: Wayfinding strategies in multi-level buildings. *Journal of Environmental Psychology* 26 (4), 284–299.
- Höök, K., 1991. An approach to a route guidance interface. Licentiate thesis, Dept. of Computer and System Sciences, Stockholm University.
- Hurtig, T., Jokinen, C., 2006. Modality fusion in a route navigation system. In: Cavedon, L., Chen, F., Dale, R., Traum, D. (Eds.), *Proceedings of Workshop on Effective Multimodal Dialogue Interfaces*. Sidney, Australia.
- Klippel, A., 2003. Wayfinding choremes. In: Kuhn, W., Worboys, M., Timpf, S. (Eds.), *Spatial Information Theory. Foundations of Geographic Information Science. International Conference COSIT*, Springer, Berlin, pp. 320–334, LNCS 2825.
- Klippel, A., Tappe, H., Habel, C., 2003. Pictorial representations of routes: Chunking route segments during comprehension. In: Freksa, C., Brauer, W., Habel, C., Wender, K. F. (Eds.), *Spatial Cognition III – Routes and Navigation, Human Memory and Learning, Spatial Representation and Spatial*

- Learning. Springer, Berlin, pp. 11–33, LNAI 2685.
- Klippel, A., Tappe, H., Kulik, L., Lee, P. U., 2005. Wayfinding choremes — a language for modeling conceptual route knowledge. *Journal of Visual Languages and Computing* 16 (4), 311–329.
- Krüger, A., Butz, A., Müller, C., Stahl, C., Wasinger, R., Steinberg, K.-E., Dirschl, A., 2004. The connected user interface: Realizing a personal situated navigation service. In: *IUI '04: Proceedings of the 9th international conference on Intelligent user interfaces*. ACM Press, New York, NY, USA, pp. 161–168.
- Kuipers, B., Tecuci, D., Stankiewicz, B. J., 2003. The skeleton in the cognitive map — a computational and empirical exploration. *Environment and Behavior* 35 (1), 81–106.
- Look, G., Shrobe, H., 2007. Towards intelligent mapping applications: A study of elements found in cognitive maps. In: *IUI '07: Proceedings of the 12th International Conference on Intelligent User Interfaces*. ACM Press, New York, NY, USA, pp. 309–312.
- Lovelace, K. L., Hegarty, M., Montello, D. R., August 1999. Elements of good route directions in familiar and unfamiliar environments. In: Freksa, C., Mark, D. M. (Eds.), *Spatial Information Theory. Cognitive and Computational Foundations of Geographic Information Science*. International Conference COSIT, Springer, Berlin, pp. 65–82, LNCS 1661.
- Lynch, K., 1960. *The Image of the City*. The MIT Press, Cambridge.
- Maaß, W., 1994. From visual perception to multimodal communication: Incremental route descriptions. *Artificial Intelligence Review Journal* 8, 68–82.
- Montello, D. R., 1998. A new framework for understanding the acquisition of spatial knowledge in large-scale environments. In: Egenhofer, M. J., Golledge, R. G. (Eds.), *Spatial and Temporal Reasoning in Geographic Information Systems*. Oxford University Press, New York, Ch. 11, pp. 143–154.
- Montello, D. R., 2005. Navigation. In: Shah, P., Miyake, A. (Eds.), *Handbook of Visuospatial Thinking*. Cambridge University Press, Cambridge, pp. 257–294.
- Patel, K., Chen, M. Y., Smith, I., Landay, J. A., 2006. Personalizing routes. In: *UIST '06: Proceedings of the 19th Annual ACM Symposium on User Interface Software and Technology*. ACM Press, New York, NY, USA, pp. 187–190.
- Porzel, R., Gurevych, I., Malaka, R., 2006. In context: Integrating domain- and situation-specific knowledge. In: Wahlster, W. (Ed.), *SmartKom - Foundations of Multimodal Dialogue Systems*. Springer, Berlin, pp. 269–284.
- Raubal, M., Winter, S., 2002. Enriching wayfinding instructions with local landmarks. In: Egenhofer, M., Mark, D. (Eds.), *Geographic Information Science - Second International Conference GIScience 2002*. Springer, Berlin, pp. 243–259, LNCS 2478.
- Richter, K.-F., 2007a. Context-specific route directions — generation of cognitively motivated wayfinding instructions. Ph.D. thesis, Universität Bremen.
- Richter, K.-F., 2007b. From turn-by-turn directions to overview information

- on the way to take. In: Gartner, G., Cartwright, W., Peterson, M. P. (Eds.), *Location Based Services and TeleCartography*. Springer, Berlin, pp. 205–214.
- Richter, K.-F., 2007c. A uniform handling of different landmark types in route directions. In: Winter, S., Duckham, M., Kulik, L., Kuipers, B. (Eds.), *Spatial Information Theory*. LNCS 4736. Springer, Berlin, pp. 373–389.
- Richter, K.-F., Klippel, A., 2005. A model for context-specific route directions. In: Freksa, C., Knauff, M., Krieg-Brückner, B., Nebel, B., Barkowsky, T. (Eds.), *Spatial Cognition IV. Reasoning, Action, Interaction: International Conference Spatial Cognition 2004*. Springer, Berlin, pp. 58–78, LNAI 3343.
- Schmid, F., Richter, K.-F., 2006. Extracting places from location data streams. In: *UbiGIS 2006 - Second International Workshop on Ubiquitous Geographical Information Services*. Münster, Germany, Workshop at GIScience 2006.
- Shanon, B., 1983. Answers to where-questions. *Discourse Processes* 6, 319–352.
- Siegel, A. W., White, S. H., 1975. The development of spatial representations of large-scale environments. In: Reese, H. (Ed.), *Advances in Child Development and Behaviour*. Academic Press, New York, pp. 9–55.
- Sorrows, M. E., Hirtle, S. C., 1999. The nature of landmarks for real and electronic spaces. In: Freksa, C., Mark, D. M. (Eds.), *Spatial Information Theory - Cognitive and Computational Foundations of Geographic Information Science*. International Conference COSIT, Springer, Berlin, pp. 37–50, LNCS 1661.
- Sperber, D., Wilson, D., 1986. *Relevance*. Basil Blackwell Ltd, Oxford, UK.
- Srivinas, S., Hirtle, S. C., 2007. Knowledge based schematization of routes. In: Barkowsky, T., Knauff, M., Ligozat, G., Montello, D. R. (Eds.), *Spatial Cognition V Reasoning, Action, Interaction*. LNAI 4387. Springer, Berlin, pp. 346–364.
- Tomko, M., 2007. Destination descriptions in urban environments. Ph.D. thesis, Department of Geomatics, University of Melbourne.
- Tomko, M., Winter, S., 2006a. Identification of the initial entity in granular route directions. In: Riedl, A., Kainz, W., Elmes, G. A. (Eds.), *Progress in Spatial Data Handling. 12th International Symposium on Spatial Data Handling*. Springer, Berlin, pp. 43–60.
- Tomko, M., Winter, S., 2006b. Recursive construction of granular route directions. *Journal of Spatial Science* 51 (1), 101–115.
- Tomko, M., Winter, S., Claramunt, C., to appear 2007. *Experiential Hierarchies of Streets*. *Computers, Environment and Urban Systems*.
- Tversky, B., Lee, P. U., 1999. Pictorial and verbal tools for conveying routes. In: Freksa, C., Mark, D. M. (Eds.), *Spatial Information Theory - Cognitive and Computational Foundations of Geographic Information Science*. International Conference COSIT, Springer, Berlin, pp. 51–64, LNCS 1661.
- Wahlster, W., Reithinger, N., Blocher, A., 2001. SmartKom: Towards multi-modal dialogues with antropomorphic interface agents. In: Wolf, G., Klein, G. (Eds.), *Proceedings of the International Status Conference "Human-Computer-Interaction"*. DLR, Berlin, pp. 23–34.

- Werner, S., Long, P., 2003. Cognition meets Le Corbusier – cognitive principles of architectural design. In: Freksa, C., Brauer, W., Habel, C., Wender, K. F. (Eds.), *Spatial Cognition III – Routes and Navigation, Human Memory and Learning, Spatial Representation and Spatial Learning*. Springer, Berlin, pp. 112–126, LNAI 2685.
- Winter, S., Tomko, M., Elias, B., Sester, M., to appear 2007. *Landmark Hierarchies in Context*. *Environment and Planning B: Planning and Design*.