A Model for Context-Specific Route Directions

Kai-Florian Richter[1], Alexander Klippel[2]

[1]Transregional Collaborative Research Center SFB/TR 8 Spatial Cognition
Universität Bremen, Germany
richter@sfbtr8.uni-bremen.de

[2]Cooperative Research Centre for Spatial Information
Department of Geomatics, University of Melbourne, Australia
aklippel@unimelb.edu.au

Abstract. Wayfinding, i.e. getting from some origin to a destination, is one of the prime everyday problems humans encounter. It has received a lot of attention in research and many (commercial) systems propose assistance in this task. We present an approach to route directions based on the idea to adapt route directions to route and environment's characteristics. The lack of such an adaptation is a major drawback of existing systems. Our approach is based on an information- and representation-theoretic analysis of routes and takes into account findings of behavioral research. The resulting systematics is the framework for the optimization process. We discuss the consequences of using an optimization process for generating route directions and outline its algorithmic realization.

1 Introduction

Getting from an origin A to a destination B is a prime problem in people's life. Efficiently solving this problem, i.e. determining a route between A and B and then purposively moving along that route, is called wayfinding (Golledge, 1999; Montello, in press). It has become a major research direction in many areas.

Wayfinding research can be organized in two broad areas: first, research that aims at shedding light on the question of how humans and other agents actually find their ways (e.g., Blades, 1991; Allen 1999); second, research that aims at supporting humans in the activity of finding a way (e.g., Wahlster et al., 2001; Heye & Timpf, 2003; Duckham & Kulik, 2003). Additionally, wayfinding can be differentiated in planning a route and actually following a route. In this contribution, we focus on supporting wayfinders in following a route.

The setting we are dealing with is wayfinding in outdoor environments where the movement occurs on a system of paths, like in city street networks or on footpaths in a park. Route directions are a primary means to guide someone in finding one's way. Here, route directions refer to instructions on how to follow a route; they are task-oriented specifications of the actions to be carried out to reach the destination (e.g., Klein, 1979; Tversky & Lee, 1998; Denis et al., 1999; Schweizer et al., 2000). Our approach complements research on incremental route directions (cf. Maaß, 1993; Habel, 2003). We use the term route directions generically in this paper to refer to any

In Christian Freksa, Markus Knauff, Bernd Krieg-Brückner, et al. (Eds.),
Spatial Cognition IV. Reasoning, Action, and Interaction:
International Conference Spatial Cognition 2004, pp. 58–78. Springer, Berlin.

form of instructions—verbal, graphical, gestures—for route following. In contrast to approaches designed to generate modality specific route directions, we present a computational model that generates *abstract route directions*, i.e. an abstract representation of the actions necessary to follow a route. The abstract representation may be *externalized* in different modalities, for example as verbal or graphical route directions or as gestures (see also Chomsky, 1986; Jackendoff, 1997; Tversky and Lee, 1999; Allen, 2003; Habel, 2003; Klippel, 2003; Klippel et al. submitted).

The article is structured as follows: we start with introducing a distinction between structure and function in wayfinding, which reflects the difference between features present in an environment and the role they take in the process of wayfinding. We argue why route directions benefit from taking into account their conceptual basis and introduce the concept of *context-specific route directions* (Section 3). A systematics of elements that can be exploited in generating abstract route directions is presented in Section 4. Section 5 motivates how this generation can be realized as an optimization process; Section 6 outlines a computational model for the optimization process and presents an example.

2 Structure and Function in Wayfinding

For the following argumentation, it is important to distinguish between the features physically present in an environment independent of any wayfinding actions and their role in the process of wayfinding. Klippel (2003) introduced the concepts of *structure* and *function* in wayfinding. With structure, he refers to an environment's physically present features; the structural level describes a static configuration of these features. Function denotes the relation of these structural elements to actions performed in the environment; the functional level demarcates those features relevant for a wayfinding action, i.e. it describes a dynamic situation and those parts of the structure that are demarcated by an action.

Accordingly, Klippel (2003, following Montello, in press) distinguishes between *path* and *route*. A path is a linear, unbounded feature in the environment upon which travel occurs. A route is a behavioral pattern; it has an origin and a destination and is directed and bounded. A route demarcates a path, i.e. it determines those parts of the paths—called *path-segments*—that are traversed while route following. We term the points where path-segments meet *branching points*. Paths, i.e. branching points and path-segments, form a *path-network*, a graph-like structure, which reflects the geometric layout of the paths in an environment, with branching points as nodes and path-segments as edges. On a functional level, we deal with *route-segments*, which correspond to those path-segments demarcated by a route. The point where two route-segments meet is termed *decision point*. At a decision point, a wayfinder needs to decide on the further direction to take; it corresponds to a branching point on the structural level. This does not imply that every decision point has to be mentioned explicitly in a route direction, as not every decision point requires the same attention by the wayfinder (see Section 4)

We consider decision points to be most pertinent for route directions. Following a route comprises two basic processes: getting to a decision point and, there, determining the further direction to take (e.g. Daniel & Denis, 1998). Route directions' main purpose is to support decision making in route following, i.e. providing information on how to proceed at a decision point. Hence, we concentrate on decision points in generating route directions; the basic representation underlying our model is a sequence of decision points.

The distinction between structure and function is also reflected in the generation of route directions (see Section 4.3). The features exploited in giving route directions are part of the structural level, i.e. they are features physically present in the environment. However, whether they are applicable for a specific route direction is determined by the functional level, i.e. in route following as well as in giving route directions the route itself demarcates those parts of an environment that are functionally relevant for the given task.

3 Conceptualizing Routes

Route directions provide instructions on how to get from A to B. Someone or some system generates them for somebody else as means of assistance; they are messages at this point. When a receiver of a message interprets or uses the instructions, they become information. Viewing route directions from this perspective allows for a representation theoretic analysis, i.e. the distinction of their syntactic, semantic, and pragmatic level (Richter et al., 2004; see also, e.g., MacEachren, 1995). The syntactic level comprises an analysis of the size of a message, for example, how many words are used in a verbal route direction, which relates to messages in an information-theoretic sense (Shannon & Weaver, 1949). On the semantic level, the processing of a message into information, i.e. the effort needed to interpret route directions, is analyzed.

Frank (2003) considers the pragmatic information content of route directions. He claims that two route directions for the same route that differ on the syntactic level, i.e. in the size of the messages, may be considered equal from a pragmatic point of view if they both lead agents to take identical routes. In this case both route directions lead to the same result, the wayfinder being at her goal, and to the same actions, the wayfinder took the same route using either of them. On the other hand, the same route direction may be different to different users, as users may differ in their knowledge of the environment or the task they try to perform with the directions given.

We are interested in the conceptual level of route directions. Two different route directions for the same route that are equal from a pragmatic perspective may well differ on the conceptual level. In order to use route directions successfully, a wayfinder needs to conceptualize the route she is about to encounter, or parts of it. As there can be differences with respect to the ease with which a route direction is understood and the extent to which it supports cognitive processes, the conceptualization of two route directions for the same route may differ. This difference may reside in the conceptualization itself. For example, the instruction 'go straight, straight, and then turn left' results in a different conceptualization than 'follow the signs to the train station'. Or the resulting conceptualizations are similar, but the differences reside in the processing of the route directions that leads to the conceptualization. For example, ‘go straight, straight, and then turn left’ may lead to the same conceptualization as ‘turn left at the third intersection’; but the former requires more processing than the latter (cf. Klippel, 2003; cf. also Dale, et al., 2002, 2003). This is because the latter directions are already chunked, while in the former directions this chunking has to be done by the wayfinder herself (see Section 4.2; see also Miller, 1956; Cowan 2001).

In our approach, *conceptualization* of route directions is the (process of forming a) mental representation of a route. A route is represented as a sequence of decision point / action pairs. Hence, more precisely, conceptualization is the (process of forming a) mental representation of an (expected) decision point sequence with their accompanying actions. We aim at creating route directions that support this conceptualization. These route directions should be easy to process, i.e. they should support forming and processing a representation of the corresponding route. Consequently, route following also becomes easier as understanding a route direction is a prerequisite for using it (cf. Dale et al., 2003).

In order to generate such route directions, we need to account for the structure of the environment in which route following takes place. The structure of an environment influences the kind of instruction that can be given. Route directions depend on the embedding of the path—instantiated by the route—in the spatial structure surrounding the path, on the structure of that path itself, on path annotations, and on landmarks that are visible along the path. Additionally, the reference system used provides alternatives to describe actions needed to follow the route. These dependencies are reflected in the systematics of elements in route directions developed in this paper. Taking into account this systematics results in abstract route directions specifically adapted to a route’s properties and the environmental characteristics. We coin the resulting route directions *context-specific route directions*[1].

This approach differs from other approaches. Duckham and Kulik (2003), for example, present an algorithm that modifies the classic AI search algorithm A*—used to find the shortest path—to calculate the route easiest to describe. They use a single way to describe a route and look for the optimal route given that description mechanism. The approach taken here starts with a given route and elicits abstract route directions for that route, which in turn are the basis for modality-adapted externalizations. Thus, our approach complements the approach of Duckham and Kulik.

---

[1] We introduce the term *context-specific route direction* to emphasize that our model explicitly adapts the generated directions to the situation at hand, i.e. to the current action to take along the route in the current surrounding environment. This reflects Dey and Abowd’s (2000) definition of context, i.e. “[…] any information that can be used to characterize the situation of an entity” (p. 3). Our model provides several alternatives to describe the same action dependent on a route’s properties and environmental characteristics. It differs, hence, from existing (internet) wayfinding assistance systems that employ strict, inflexible rules to any context. Such strict rules lead to effects like leaving a city when entering an inner-city highway and getting back to the same city when exiting that highway again, or the generation of new events just because the name of a street changes, not because the wayfinder needs to change her current action.

It also differs from the CORAL system by Dale et al. (2002, 2003). Contrary to their approach we are not restricted to natural language output; the 'best' route direction also could mean a graphic representation or some mixed modality like map gestures (Hirtle, 2000) of a route or a part thereof. Common aspects are discussed, for example, in Section 4.2. The conceptualization approach taken here aims at an abstract representation formalism that forms the basis for various output formats or different realizations of the same output format.

Guhe et al. (2003) present an approach to generate abstract representations of motion events, which can be extended to the generation of route directions (Habel, 2003). They aim for a system that is able to describe events in a dynamically changing world. There are two main differences to our approach: first, as with the approach by Dale et al., their model is intended for natural language output; second, they focus on processing dynamic situations where information is acquired incrementally, and accordingly, their interest is in incremental route instructions.

4 A Systematics for Context-Specific Route Directions

The structure of an environment and its elements need to be considered in creating abstract route directions (ARDs) since it influences how instructions for route following can be given. This influence can be local, i.e. an environment's element is usable for one or a few ARDs, or global, i.e. an environment's structure is exploitable for several or most of the ARDs. Furthermore, the reference system used determines alternatives to give instructions. An action needed for route following can be described from the perspective of the wayfinder (egocentric references). Elements of the environment can be referred to in instructions (allocentric references), or some fixed references outside the environment may be used (absolute references). In these references, different elements of an environment and of a route may be employed for giving route directions. The elements are cataloged according to three levels in our systematics.

4.1 Levels of the Systematics

The three levels of the systematics reflect three categories of elements that can be used in giving route directions: global references, i.e. elements that are not part of the immediate surrounding, environmental structure, i.e. elements of an environment that impose a structure on that environment, and elements that belong to the path and the route.

Global References

This level comprises elements referred to in abstract route directions that are not part of the immediate surrounding environment in which the action takes place. These are references that rely on an absolute reference system, i.e. the direction referred to is the same everywhere in the environment and does not depend on a wayfinder's position. Most typical are cardinal directions like 'north', 'east', etc. Additionally, references to global landmarks, i.e. landmarks outside the surrounding environment, belong to this category if references to them are the same everywhere in the given environment. These landmarks are visible from many places of the environment or their location is everywhere unequivocally known, which makes them usable as reference objects (Sorrows & Hirtle, 1999). An example for a resulting instruction is ‘towards the sea’ with ‘the sea’ being an example of such a landmark.

Environmental Structure

There can be elements that are part of an environment, which have an influence on an environment as a whole. Such elements impose a structure on that environment, which leads to distinctive parts. Hence, these elements or a reference to the emerging distinctive parts can be exploited in abstract route directions. Since the parts are distinctive, they provide unambiguous direction information usable in giving route directions. An example for such an element is a slant; examples for such instructions are ‘uphill’ or ‘downhill’.

Paths, Routes, and Landmarks

The third level comprises elements that relate to paths and routes. Landmarks and decision points are part of this level, as well as *annotations* along a path, which are elements that are set up to unequivocally identify that path, like street names, street signs, or markers. We also catalog the combination of several instructions into a single one to be on this level (Dale et al., 2003; Klippel et al., 2003).

Typically, instructions on this level use an egocentric reference frame, like ‘turn right’ or ‘go straight’. That is, route directions refer to the locations of an intersection’s branches relative to the wayfinder. We need to consider the configuration of a branching point when creating ARDs on this level. Since our aim is to generate ARDs that are unambiguous, the branch to take when following the route needs to be unequivocally identifiable. From a functional perspective, several branches at a decision point may broadly lead to the same direction and, therefore, the functionally relevant branch needs to be further specified. For example, if there are two branches leading to the left, they may be distinguished into ‘half left’[2] and ‘sharp left’ and the resulting instruction may be ‘turn half left’.

Landmarks are pertinent to route directions (cf., e.g., Denis et al., 1999). They influence the way instructions are given. For example, people refer more often to landmarks than to street names in generating route directions and they are more effective than street names in route guidance (cf. Tom & Denis, 2003). Structurally, landmarks can be point-like, linear, or areal. Point-like landmarks are located in small, restricted areas of an environment. Such a landmark is, for example, a salient building like a church. The other two, linear and areal landmarks, extend across an environment, like a river or a forest. We consider landmarks that influence route directions to be part of the route. We call them *routemarks* (cf. Krieg-Brückner et al, 1998). A routemark can be either at a decision point, at a route-segment between two decision points (Hermann et al., 1998), or in some distant to, but visible from the

---

[2] The German term *halb* [*links* / *rechts*] is not translatable directly to English. More appropriate natural language expressions would be *veer*, *bear*, *up to the* [*left* / *right*] etc.

route. We call the latter kind of routemark *distant routemarks* (cf. Lovelace et al., 1999).

Functionally, routemarks at decision points can be used to identify a decision point, for example, ‘turn left *at* the church’. Routemarks between decision points may be employed to further describe the route and to function as confirmation that one is still on the right track (‘you *pass* a church’). Distant routemarks, finally, are like beacons. Assuming they are visible while passing several decision points, they can be used as pointers to a certain direction. An example of such an instruction is ‘*towards* the church’. For the conceptualization of a turning action, the location of a landmark at a decision point is important. A routemark may be passed before the turn (“turn *after* the church”), after the turn (“turn *before* the church”), or the landmark may not be located at a functionally relevant branch of the decision point (“turn where the church is”) (cf. Klippel, 2003). Routemarks before a turn are easily conceptualized as the turning action occurs immediately after them. They are, thus, a good identifier for a decision point. It is an open issue, though, what influence the latter two cases have on conceptualization and which additional parameters play a role here. We introduce and distinguish them for reasons of completeness in the systematics.

On a functional level, linear and areal landmarks can function either point-like or linear. In a point-like fashion, such landmarks identify a decision point and may indicate an action to be taken, for example ‘turn right when coming to the river’. However, linear and areal landmarks not only identify a decision point, but may also allow combining several decisions into one decision. Examples for route directions that usually involve a linear pattern are ‘follow the river’ or ‘walk along the forest’. They may determine the actions for several decision points. Therefore, such route directions require an additional qualifier that establishes the point until the instruction holds. An example for such a qualifier is ‘until you reach the gas station’.

4.2 Chunking Instructions

Route directions provide instructions on how to proceed for every decision point. Yet, not every decision point and the accompanying action need to be mentioned explicitly. Often, it is possible to combine actions for several decision points into one route direction; this combination is an important mechanism in route directions and the conceptualization of routes. We call it *spatial chunking* (Klippel et al., 2003). Dale et al. (2003) refer to it as *segmentation.*

Spatial chunking groups several decision point / action pairs into a single segment; we call these segments *higher order route direction elements* (HORDE) (cf. Klippel, 2003). Dale et al. (2003) identify two segmentation principles: *landmark-based* and *path-based* segmentation. In landmark-based segmentation, landmarks at decision points delimit a part of the route to be followed; the route is decomposed into segments, each leading to such a landmark. Path-based segmentation is based on three features of paths—road status (highways, main roads, etc.), path length, and turn saliency (e.g. T-intersections). By employing any of these features or a combination thereof, routes can be segmented.

Klippel et al. (2003) differentiate three kinds of spatial chunking:

- Numerical chunking: Here, a sequence of several decision points that involve no direction change (termed *DP-*) and one decision point with a direction change (*DP+*) are combined into a single decision. This is done by counting the decision points until a direction change occurs, for example 'turn left at the third intersection'. Also, a sequence of decision points with equal direction changes can be grouped, for example 'turn twice right'.
- Landmark chunking: This kind of chunking is similar to numerical chunking. However, an unambiguous landmark identifying the DP+ is utilized to mark the point where a direction change occurs, instead of counting the DP-. An example for such a HORDE is 'turn right at the gas station'. The number of intermediate decision points is not specified in this kind of chunking.
- Structure chunking: In structure chunking, spatial structures that are unique in a given local environment are exploited. For example, the dead end of a T-intersection unequivocally marks the need for a direction change—one either needs to turn left or right as straight on is impossible. Hence, it is possible to chunk several DP- and the relevant DP+ located at such a structure into HORDE like 'turn right at the T-intersection'. An instruction like 'follow the river' also rests upon structure chunking as it combines actions for several decision points that are located along the river into a single one.

Klippel et al. (2003) present spatial chunking based on an egocentric reference frame and on instructions employing elements from the third level of the systematics used here. Abstract route directions on the other levels of the systematics can also be chunked, i.e. it is possible to combine sequences of several decision points with ARDs that employ other elements of the systematics. To pick up the example from the previous subsection again, the instruction 'go uphill' may combine actions for several decision points that happen to be in a line uphill into a single decision, which can also considered to be a HORDE. Usually, either landmark or structure chunking are used to combine ARD for several decision point / action pairs on the higher levels of the systematics. Just like with linear landmarks, there needs to be a qualifier that marks the end of such a HORDE, i.e., that denotes the point until an instruction based on such a HORDE holds.

4.3 Structure and Function in the Systematics

The distinction made between structure and function in wayfinding is also reflected in our systematics. The elements presented above are all part of the structural level, i.e. they are all part of either the path itself or the environment the path is embedded in—with an exception of global landmarks, which are nonetheless also clearly part of the structural level as references to them do not depend on a wayfinder's location in the environment. However, whether these elements are applicable in creating route directions for a route depends on the context set by the functional level.

The route, which is on the functional level, demarcates the functionally relevant parts of an environment, i.e. the path segments traversed while route following. It also determines the actions needed to follow that route; the corresponding sequence of decision point / action pairs represents the dynamic aspects of route following, especially the direction in which a wayfinder moves through the environment.

Abstract route directions need to reflect this sequence and provide information on the directions to take. Consequently, only elements of the systematics that unequivocally denote these directions, i.e. allow a wayfinder to correctly orient herself, are applicable in generating ARDs.

4.4 Granularity in the Systematics

Granularity is one of the fundamental aspects of knowledge representation (cf. Hobbs, 1985). The elements of the systematics, which offer access to knowledge of an environment, are on different levels of granularity, i.e. ARDs generated using these elements provide information on how to follow a route on different levels of granularity. These changes in granularity reside between the levels of the systematics as well as within the levels. In our systematics, granularity refers to how closely abstract route directions are linked to individual decision point / action pairs and the corresponding branching point's configuration, i.e. to what extent ARDs abstract from a detailed description of a decision point / action pair itself.

The level of global references, which is the first level of the systematics, relates to the coarsest granularity. Referring to elements that are not part of the surrounding environment results in coarse direction information, which is not explicitly based on the structure of the environment itself. Such instructions just exploit that they lead to unequivocal choices at the decision points they hold for. ARDs on the level of environmental structure—the second level of the systematics—provide still coarse information on the further direction to take, but which explicitly takes into account an environment's structure and is, therefore, more closely related to the embedded path itself. Consequently, this kind of instruction is on a finer granularity level than those using elements of the systematics' first level.

The third level, which is the level of path, route, and landmarks, contains elements of the route itself. ARDs generated with these elements usually refer explicitly to decision points. The structure of these ARDs is therefore close to the decision point / action pairs themselves. Accordingly, we consider elements on the third level to be on the finest level of granularity.

A change of granularity occurs also within the different levels of the systematics. Chunking instructions obviously increases the degree of abstraction from individual decision point / action pairs. This holds for all three levels. The different kinds of chunking result in abstract route directions on different levels of granularity. While landmark and structure chunking combine a number of decision points that are not specified in the resulting instruction, i.e. they abstract from the exact number of decision points involved and may therefore provide a single instruction for a large part of the route, in numerical chunking the number of decision points is explicitly mentioned. Such instructions are only sensibly applicable for a small number of decision points (cf. Klippel, 2003). Finally, taking into account the configuration of a branching point, for example, by further qualifying a turning instruction is on the finest level of granularity as this is directly based on an individual decision point.

4.5 Implicit vs. Explicit Representation

We use a sequence that contains every decision point of a route as an underlying representation in our model. The application of chunking, however, combines several of these decision points into a single representation. Thus, the resulting representation does not necessarily contain every single decision point anymore. However, when following a route, a wayfinder needs to make a decision at every decision point encountered along the route. Therefore, she must be able to infer the decisions only implicitly represented in the route directions in order to know the further direction to take. Consequently, the route directions need to be correct, i.e. provide instructions on a route that leads from origin to destination, and complete, i.e. provide the instructions such that every decision necessary can be derived from them.

5 Generating Context-Specific Route Directions: An Optimization Problem

Our aim is to generate abstract route directions, which form the basis for real route directions that ease the conceptualization of a route. To this end, we need to choose from all possible ways to create abstract route directions for a route the one that best fits this aim. More precisely, for each decision point along the route we need to choose an (abstract) instruction on which action to perform that is most likely to ease the conceptualization. However, the kind of instruction to choose at a decision point may depend on the kind of instruction chosen for previous or following decision points, as, for example, decisions may be chunked. Thus, all different kinds of abstract route directions that are possible for a decision point need to be judged according to their consequences regarding conceptualization, taking into account the possible abstract route directions for other decision points of that route. The dependence of a local choice (an instruction for a single decision point) on the choices made elsewhere (instructions for other decision points) clearly shows that we are dealing with an optimization problem; we are looking for optimized abstract route directions for a given route. Accordingly, the sequence of decision point / action pairs needs to be processed and translated into optimized abstract route directions for each pair.

The kinds of abstract route directions that can be created for a decision point / action pair are based on the systematics presented in the last section. For the generation of abstract route directions, we explicitly exploit an environment and route's characteristics; this results in ARDs specifically adapted to these characteristics.

If we want to decide on which ARD for a decision point / action pair best fits our aim of an easy conceptualization, we need a measure to compare possible ARDs with respect to that aim. We need rules that define which kind of ARD to choose in which situation, i.e. an optimization criterion. Potentially, there is a huge number of such rules. There can be many kinds of abstract route directions applicable at the same time for a decision point, which need to be judged. Moreover, in combination with the potential dependency on ARDs chosen for other decision points, the number of rules needed increases even further. Thus, computationally, it is not sensible to have a specific rule for every situation that might occur in creating context-specific route directions. Instead, we need a heuristic, i.e. general rules that provide guidelines and sensible choices that can be applied when a specific situation occurs. As in most optimization problems, these general rules may in some cases result in abstract route directions that are not the best possible—though still good ones; but applying these heuristics makes the problem computationally feasible.

Several optimization criteria may be applicable. A first simple heuristic would be to use always the highest granularity level possible, i.e. to choose for each decision point the ARD that corresponds to an element of the systematics, which is—compared to all other possible elements—on the highest granularity level in the systematics. Other possible criteria include: minimal number of distinct parts, i.e. smallest number of chunks; abstract route directions based on no more than n elements of the systematics, i.e. the optimization results in abstract route directions that do not employ more than n different elements of the systematics; no more than n changes in the kind of instructions, i.e. in the resulting abstract route directions there is at most n times a switch from one element of the systematics to another; no more than n changes in the reference system used.

In order to create context-specific route directions we propose to aim at a minimal number of distinct parts with, everything else being equal, abstract route directions on the highest granularity levels possible. In the following, we will argue for this criterion.

First, from an information-theoretic perspective a small number of chunks reduces the amount of information that needs to be communicated, i.e. the size of the message decreases. A decrease of information leads to a decrease of memory load, i.e. the wayfinder needs to remember less information. To put it another way, a reduction in number of chunks results in a decreased amount of information explicitly represented and an increase of information that needs to be inferred. This relates to Grice's (1975) principles of communication, especially the ones he termed *quality* and *quantity*: the information provided needs to be correct and should not contain any details that are unnecessary for the message's purpose.

Second, the application of chunking and HORDE also reduces the processing involved in conceptualizing a route. As argued before, a principle of cognitive ergonomics is to combine several instructions into a higher-order instruction if possible. This clearly requires additional processing of the route directions, i.e. increases the cognitive load of a wayfinder. Since in the generation of context-specific route directions this chunking of single ARDs to higher-order route directions is already done, the wayfinder does not need to perform this herself anymore, which, accordingly, eases the cognitive processing of the route directions. That is, with context-specific route directions we provide instructions for route following that are easy to process and theoretically easy to memorize.

Furthermore, route directions on a high level of granularity reduce the problem of matching an expected decision point / action pair with the real environment. This kind of instruction is less prone to errors if the conceptualized decision point / action pair does not (exactly) match the actual situation in the environment. For example, an instruction 'turn left' might get a wayfinder into trouble if the actual configuration of branches met at an intersection does not seem to include a branch she considers leading to the left. While an instruction ‘follow the signs to the train station’ does not depend at all on the configuration of the intersections; all that is required is that there is actually a sign pointing in direction to the train station. Thus, with route directions on higher levels of granularity a wayfinder is not that strongly dependent on the environment meeting her conceptualization anymore.

It can also be argued that applying HORDE and providing route directions on a high level of granularity moves the task of wayfinding in an environment nearer to the task of planning a trip through an environment. Such route directions include fewer “real” decision points, i.e. fewer decision points where a wayfinder actively needs to remember a direction change (DP+). As HORDE combine several decision points into one decision, a wayfinder only needs to remember the point until the HORDE holds; all decisions in between can be inferred. ‘Turn right at the third intersection’, for example, indicates a direction change at the third intersection a wayfinder encounters; the information implicitly represented is that she has to keep the current direction at the first and second intersections. The advantages of HORDE are even more obvious when looking at instructions like ‘follow the markers’ for a hiking trail. Here, a single instruction suffices to lead a wayfinder to her goal; but following the markers may involve many direction changes while walking along the hiking trail. That is, such a HORDE on a high level of granularity may render decision points that actually involve a direction change, i.e. DP+, into decision points that do not require a change of action, i.e. practically turn them into DP-.

6 A Computational Model for Context-Specific Route Directions

The generation of context-specific route directions (CSRD) can be realized as an optimization problem. We need to find globally optimal abstract route directions for each decision point / action pair or chunks thereof, respectively, i.e. ARDs that are optimal with respect to the complete route, not just for a single decision point / action pair. In the last section, we presented the optimization criterion employed in our approach. In this section, we provide an overview on the computational part of our approach, which includes the algorithm used for finding the optimal CSRD. Additionally, we give an example of how the optimization process works and discuss how to deal with missing data.

6.1 A Computational Approach to Context-Specific Route Directions

For the automatic generation of CSRD we need information on the route in question, i.e. we need a representation of the environment that contains all information needed and allows us to compute a route from some origin to a destination. To this end, we employ a graph-like representation of the environment’s path-network. The graph’s edges represent the path-segments; nodes denote the branching points. The graph reflects the layout of the environment’s paths, i.e. it preserves information on angles between branches and distances. In such a graph, we can calculate a route with any path-search method, like A* or Dijkstra’s (1959) shortest path algorithm. The calculation results in a sequence of nodes that need to be traversed to get from an origin to a destination. This sequence corresponds to the sequence of decision points that is the underlying representation of our model. For the generation of context-specific route directions we need additional information on the elements of the systematics, for example on position, structure, and visibility of landmarks or on path annotations. Hence, we annotate the graph with this information (see Section 6.3 for a discussion on automatically extracting such information).

For the optimization process, we start with generating for each decision point all abstract route directions (ARDs) that are possible according to the systematics defined. Such ARDs represent a decision point and its accompanying action description based on the element used. The action description consists of a direction relation and, if one applies, that feature of the environment the relation refers to. Examples of ARDs are (DP$_1$,`left`), denoting a left turn at the first decision point of a route, or (DP$_4$,`follow/marker`), representing an instruction to follow the marker at the fourth decision point. In case of a routemark at a decision point, we also employ a relation to denote the position of that landmark: `after` is used as a relation to state that a turn occurs after a landmark is passed; `before` to state that a turn occurs before a landmark is passed, and the relation `at` is a generic term representing the presence of a landmark anywhere at a decision point (see Section 4.1). Thus, (DP$_2$,`right/after church`) denotes a right turn that can be further qualified using a landmark, here a church, which is passed before the turn occurs.

Each element of the systematics has a corresponding set of direction relations; these differ across the elements. For ARDs based on egocentric references, for example, we use the relations defined in the sector model presented in Klippel (2003), which has been further refined in behavioral experiments (e.g., Klippel et al., 2004). The model comprises three basic directions—`straight`, `left`, `right`—and two additional qualifiers for left and right—`half` and `sharp`—leading to seven different directions. As another example, global references are either represented with a cardinal direction—`north`, `east`, `south`, `west`—or with the relation `towards` combined with a referenced global landmark, like `towards/sea`.

The relations used in this approach, like `left` or `towards`, represent information on the direction to take at a decision point. This resembles the symbolic operators describing directional phrases as, for example, in Jackendoff (1990) or Eschenbach et al. (2000). However, it is important to note that all relations used and the abstract route directions, like (DP$_2$,`right/after church`), are by no means meant to be the actual (verbal) output of an assistance system. They are an abstract representation of the systematics' elements applicable for a given decision point, i.e. they represent possibilities of how a decision point / action pair can be described according to the systematics. We choose relation terms like `left` or `towards` because they are more readable than terms like `a`, `b`, `c`, and so on, but these terms need not be the terms used in an actual verbal output. The step to generate verbal or graphical route directions presented to a user, i.e. the transformation of the abstract route directions into concrete ones, is not covered in this research.

We check for each decision point which elements of the systematics are applicable and generate an abstract route direction based on this element. The annotations in the street-network's graph provide information on which elements can be used, for example, whether a landmark is located at a decision point or whether a global landmark is visible. This way, possible ARDs are generated resulting in a set of instructions for every decision point of the route.

Our aim is to find a minimal number of distinct parts in the abstract route directions on the highest granularity levels, i.e. in our route directions, we try to cover the complete route with as few chunks as possible while using elements of the systematics on the highest possible granularity levels. This resembles the approach of Dale et al. (2002): "... the general idea is to view messages as data objects corresponding to the largest distinct linguistic fragments we need in order to generate the variety of texts we are interested in" (p. 4). Different to Dale et al., our chunks are not necessarily “linguistic fragments” found in natural language route directions, but are derived from spatial data according to principles of HORDE.

We are looking for sub-sequences in the decision point sequence that share abstract route directions based on the same elements of the systematics and are chunkable. We apply the chunking rules as described in Section 4.2; these can be further refined to exclude results of the chunking process that are not sensible. Klippel (2003), for example, derived a set of rules in his wayfinding choreme route grammar for generating valid HORDE based on the direction model explained above. Other rules, like those by Dale et al. (2003) or the route direction principles by Denis (1997), can also be incorporated in our optimization process to prevent insensible chunks like ’right at the 21st intersection’.

For the optimization process, we choose the first ARD of the first decision point and calculate the union with the following decision points’ sets of ARD until we encounter a decision point, which cannot be chunked with the previous ones according to the chunking rules employed. We then choose the next ARD of the first decision point and again try to chunk it with as many of the following decision points as possible. We repeat this until all abstract route directions of the first decision point have been processed. We continue with the second decision point, again building chunks with every possible ARD for that decision point. The process runs until we generated all chunks for every decision point of the route. Along this process, we keep track on which combination of chunks is minimal, i.e. covers the most decision points with the least number of chunks. The process can be implemented using dynamic programming; Table 1 summarizes it.

6.2 An Example

To clarify the idea of optimization, we present an example of our approach. We chose the Bürgerpark in Bremen—a big park in the center of the city. Route directions are generated from one of its entrances to one of the park’s cafés. Fig. 1 shows a schematic map of the area; the chosen route is shown as a black line.

The network of path segments is represented as a graph (see Fig. 2). The edges that correspond to the path-segments demarcated by the chosen route are shown as a bold line. The graph is also annotated with information on landmarks, like the buildings shown in Fig. 1; the annotations are not shown in Fig. 2. The route consists of ten decision points, i.e. ten branching points are passed when following this route.

Table 1. The optimization process in an algorithmic description.

```
For each DP in route,
  determine every ARD possible according to the
  systematics resulting in a set of ARD.

Start with first ARD of first DP,
  try to chunk it with as many following DPs as
  possible,
  store generated chunk.

Repeat with following ARDs of first DP,
  until all ARDs have been processed.

Store biggest chunk as current CSRD.

Repeat with following DPs.
  If,
    the biggest newly generated chunk does not overlap
    with current CSRD, add to CSRD.
  else if,
    the newly generated chunk covers more DPs than the
    one covering the current DP currently stored in
    CSRD, rebuild CSRD using new chunk.
  until all DPs have been processed.
```

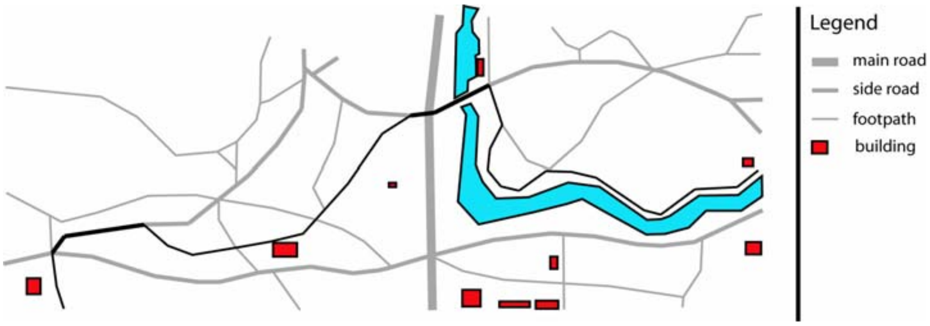

Fig. 1. A schematic map showing a part of the Bürgerpark in Bremen.

As a first step in the optimization process, we determine for each decision point all ARDs that are possible according to the systematics (see Table 2). According to the sector model used (cf. Section 6.1; Klippel et. al, 2004), there is a half-right turn at the first decision point. There is no further information available on this decision point in the graph, hence (`DP`$_1$`,half right`) is the only abstract route direction that can be generated for this decision point. The same holds for the second decision point. At decision point three, there is no change in direction; the direction relation applicable is `straight`. Here, additionally a landmark—a building in the park—visible from this decision point is in the direction of movement. It functions like a beacon as described in Section 4. This can be exploited resulting in two possible abstract route directions for this decision point: (DP$_3$,`straight`), (DP$_3$,`towards/building`). We continue this process until all possible ARDs for all decision points have been generated; these are shown in Table 2. Decision points six, nine, and ten are worth a closer examination: the direction change at decision point six is slightly to the right (`half right`). However, since this turn is at a T-intersection, we can exploit this structural element rather than relying on the direction concept alone. The direction relation is coarsened to `right`—since at a T-intersection only a left or right turn is possible—and the structural element is added resulting in (DP$_6$,`right/T-intersection`). For decision points nine and ten, again `half right` is the relation to use according to the employed sector model; but here also a linear landmark—the river—is exploitable, as the route segment traveled resides along the river. There are, thus, again two possible abstract route directions: (DP$_9$,`half right`) and (DP$_9$,`follow/river`) (the same for decision point ten).

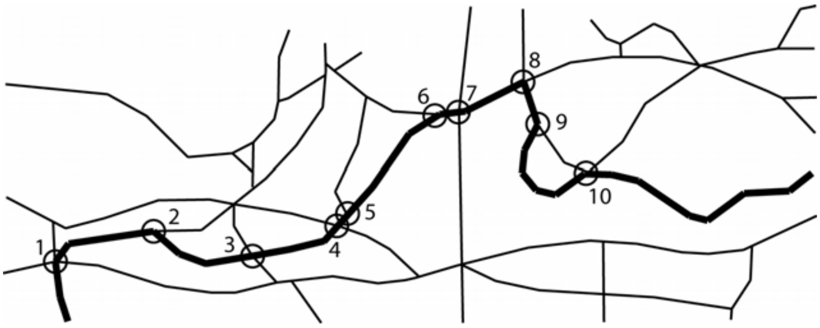

Fig. 2. The graph corresponding to the network of ways (see Fig. 1). The dots mark the decision points, which are numbered in the order of passage.

For the optimization process, we choose the first ARD in the set of the first decision point and try to chunk as many decision points as possible applying chunking rules for the kind of ARD chosen. Here, we can apply numerical chunking for the first two decision points, i.e. chunk both abstract route directions `half right`. This chunk is also the current optimal CSRD—since it is the only chunk so far. As there are no more ARDs for the first decision point, we continue with the second decision point. Its ARD cannot be chunked with any of the third decision point's ARDs. Therefore, we store this single abstract route direction for the second decision point; the CSRD still consists of the chunk generated for the first decision point.

Table 2. The route's decision points and their set of possible abstract route directions[3].

Decision point	Set of abstract route directions
1	`{half right}`
2	`{half right}`
3	`{straight, towards/building}`
4	`{straight}`
5	`{straight, towards/building}`
6	`{half right, right/T-intersection}`
7	`{straight, straight/at bridge}`
8	`{right, right/after bridge, right/at building}`
9	`{half right, follow/river}`
10	`{half right, follow/river}`

Decision point three to six can be chunked using structure chunking (three times `straight`, followed by `right/T-intersection`). This is the biggest chunk that can be generated for these four decision points. Our CSRD now consists of two chunks; the first grouping decision points one and two, the second grouping decision points three to six. Finally, for the last four decision points, the best chunks we can generate are for seventh and eighth decision point (`straight/at bridge`, `right/after bridge`) and for ninth and tenth (`follow/river`, `follow/river`). Thus, the abstract route directions for the resulting CSRD consist of four chunks; they are summarized in Table 3. They contain the chunked decision points in the first part, and the direction relations that get chunked in the second part.

Table 3. The resulting chunks of decision points for the chosen route.

{DP_1,DP_2;half right,half right}
{DP_3,DP_4,DP_5,DP_6;straight,straight,straight,right/T-intersection}
{DP_7,DP_8;straight/at bridge,right/after bridge}
{DP_9,DP_{10};follow/river,follow/river}

6.3 Availability of Data

The success of our approach relies on the availability of data on the environment route following takes place in. The underlying representation of our model is a sequence of decision points passed along a route. The graph used to calculate this sequence is derived, most typically, from GIS data sets like the ones provided by federal authorities. For the purpose of creating context-specific route directions, this graph needs to be further annotated with additional information, like street signs or landmarks.

---

[3] We omitted cardinal directions in this example to keep it simple. No chunks based on cardinal directions would contribute to the resulting route directions; thus, this omission does not change the presented optimization process.

While the availability of land-use data is fairly good these days and, thus, such a graph is readily available, additional data is not systematically available. Still, we argue that our approach calculates reasonable results even with missing data. In recent years, spatial, especially geographic data has increased tremendously in its importance for business; accordingly, more and more data gets collected. The automatic extraction of such data has become an important research issue. For example, there is work on extracting landmarks from land-use data and on automatically determining the saliency of landmarks (cf. Elias & Sester, 2003; Winter, 2003; Raubal & Winter, 2002, respectively).

Most importantly, even if such additional data is not completely available our approach still achieves good results. The optimization process as described above makes use of whatever data is available. It optimizes route directions according to this data. If only the underlying graph of the path network should be available, it is still possible to create route directions using an egocentric reference system, i.e. using directions like 'turn left' or 'go straight', and to apply spatial chunking to combine these into HORDE. Consequently, this still results in route directions, which are as good or even better than those generated by assistance systems available today.

7 Conclusions & Outlook

We present an approach to generate abstract route directions that explicitly takes into account a route's properties and environmental characteristics. This is a first step to *context-specific route directions*. They support the conceptualization of a route as they reflect cognitive principles of organizing spatial knowledge. To generate such route directions automatically, we employ an optimization process. This process aims at minimizing the number of distinct parts of route directions. Our model is based on a systematics of elements that can be employed in creating abstract route directions; this systematics reflects different levels of granularity and respects the distinction between structure and function in wayfinding.

Our claim is that context-specific route directions are easier to conceptualize, i.e. they allow forming a mental representation of a route that is easier to process and that better matches the actual route encountered. Hence, route following becomes easier. Our approach adapts abstract route directions to actual situations in the environment. Compared to existing approaches that use the same references and set of actions irrespective of the route, which may lead to inadequate, hard to use route directions, this is, thus, a step towards the goal of providing context-specific route directions that support cognitive processes (cf., e.g., Dale et al., 2003, for a critique on existing internet route-planners and Habel, 2003, for a discussion of benefits of multimodal route instructions). The approach differs from those that aim to specify natural language processes in that the scope of *conceptualization* is extended to information available in spatial data and conceptualization processes that are beyond those required for natural language generation.

Future work comprises an extension of the presented systematics. Furthermore, we need to evaluate different optimization criteria, applying both behavioral research and computational specification, to refine our approach. Finally, with some adaptation, our approach may also be usable to calculate routes through an environment that are optimized with respect to their ease of conceptualization, i.e. to already account for the proposed optimization in the path-search algorithm. This is in line with approaches like Duckham and Kulik's (2003), which try to overcome today's wayfinding assistance systems' limitations of just calculating shortest or fastest routes.

Acknowledgements

This research is carried out as a part of the Transregional Collaborative Research Center SFB/TR 8 Spatial Cognition. Funding by the Deutsche Forschungsgemeinschaft (DFG) is gratefully acknowledged. We like to thank Christian Freksa and Heike Tappe for helpful comments and discussions. We also thank three anonymous reviewers.

References

Allen, G.L. (1999). Spatial abilities, cognitive maps, and wayfinding. Bases for individual differences in spatial cognition and behavior. In R.G. Golledge (ed.), *Wayfinding behavior. Cognitive mapping and other spatial processes* (pp. 46-80). Baltimore: John Hopkins University Press.

Allen, G. L. (2003). Gestures accompanying verbal route directions: Do they point to a new avenue for examining spatial representations? *Spatial Cognition and Computation* 3:259-268.

Blades, M. (1991). Wayfinding theory and research: The need for a new approach. In D.M. Mark & A.U. Frank (eds.), *Cognitive and linguistic aspects of geographic space* (pp. 137-165). Dordrecht: Kluwer.

Chomsky, N. (1986). *Knowledge of language: Its nature, origin and use.* New York: Praeger.

Cowan, N. (2001). The magical number 4 in short term memory. A reconsideration of storage capacity. *Behavioral and Brain Sciences* 24:87-186.

Dale, R., Geldof, S., and Prost, J.-P. (2002). Generating more natural route descriptions. In *Proceedings of the 2002 Australasian natural language processing workshop.* Canberra, Australia.

Dale, R., Geldof, S., and Prost, J.-P. (2003). CORAL: Using natural language generation for navigational assistance. In *Proceedings of the 26th Australasian computer science conference (ACSC2003).* Adelaide, Australia.

Daniel, M-P. & Denis, M. (1998). Spatial descriptions as navigational aids: A cognitive analysis of route directions. *Kognitionswissenschaft* 7(1):45-52.

Denis, M. (1997). The description of routes: A cognitive approach to the production of spatial discourse. In *Cahiers de Psychologie Cognitive* 16(4):409-458.

Denis, M., Pazzaglia, F., Cornoldi, C., and Bertolo, L. (1999). Spatial discourse and navigation: An analysis of route directions in the city of Venice. *Applied Cognitive Psychology* 13(2):145-174.

Dey, A.K. & Abowd, G.D. (2000). Towards a better understanding of context and context-awareness. In *Proceedings of the workshop on the what, who, where, when, and how of context-awareness*. New York.

Dijkstra, E.W. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik* 1:269-271.

Duckham, M. & Kulik, L. (2003). Simplest paths: Automated route selection for navigation. In W. Kuhn, M. Worboys, and S. Timpf (eds.), *COSIT 2003. Spatial information theory* (pp. 169-185). Berlin: Springer.

Elias, B. & Sester, M. (2003). Landmarken für Wegbeschreibungen - Identifikation, Extraktion und Visualisierung. *Kartographische Nachrichten* 53(2):51-57.

Eschenbach, C., Tschander, L., Habel, C., and Kulik, L. (2000). Lexical specifications of paths. In C. Freksa, W. Brauer, C. Habel, and K.F. Wender (eds.), *Spatial Cognition II. Integrating abstract theories, empirical studies, formal methods, and practical applications* (pp. 127-144). Berlin: Springer.

Frank, A.U. (2003). Pragmatic information content. How to measure the information in a route description. In M. Goodchild, M. Duckham, and M.Worboys (eds.), *Perspectives on geographic information science* (pp. 47-68). London: Taylor and Francis.

Golledge, R.G. (1999). *Wayfinding behavior. Cognitive mapping and other spatial processes*. Baltimore: John Hopkins University Press.

Grice, H.P. (1975). Logic and conversation. In P. Cole & J.L Morgan (eds.), *Speech acts* (pp. 41-58). Volume 3 of *Syntax and semantics*. New York: Academic Press.

Guhe, M., Habel, C., and Tschander, L. (2003). Describing motion events: Incremental representations for incremental processing. In M. Guhe, C. Habel, and L. Tschander (eds.), *Proceedings of the 5th international workshop on computational semantics (IWCS-5)* (pp. 410-424). Tillburg.

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.

Heye, C. & Timpf, S. (2003). Factors influencing the physical complexity of routes in public transportation networks. *Electronic proceedings of the 10th international conference on travel behaviour research*. Lucerne.

Herrmann, T., Schweizer, K., Janzen, G., and Katz, S. (1998). Routen- und Überblickswissen - konzeptuelle Überlegungen. *Kognitionswissenschaft* 7(4):145-159.

Hirtle, S.C. (2000). The use of maps, images and "gestures" for navigation. In C. Freksa, W. Brauer, C. Habel, and K. F. Wender (eds.), *Spatial Cognition II – Integrating abstract theories, empirical studies, formal methods, and practical applications* (pp. 31-40). Berlin: Springer.

Hobbs, J.R. (1985). Granularity. In A.K. Joshi (ed.), *Proceedings of the 9th international joint conference on artificial intelligence* (pp. 432-435). San Francisco: Morgan Kaufmann

Jackendoff, R. (1990). *Semantic structures.* Cambridge, MA: MIT Press.

Jackendoff, R. (1997). *The architecture of the language faculty.* Cambridge, MA: MIT Press.

Klein, W. (1979). Wegauskünfte. *Zeitschrift für Literaturwissenschaft und Linguistik* 9(33):9-57.

Klippel, A. (2003): *Wayfinding Choremes – Conceptualizing wayfinding and route direction elements*. Universität Bremen.

Klippel, A., Tappe H., and Habel, C. (2003). Pictorial representations of routes: Chunking route segments during comprehension. In C. Freksa, W. Brauer, C. Habel, and K.F. Wender (eds.), *Spatial Cognition III. Routes and navigation, human memory and learning, spatial representation and spatial learning* (pp. 11-33). Berlin: Springer.

Klippel, A., Dewey, C., Knauff, M., Richter, K.-F., Montello, D. R., Freksa, C., and Loeliger, E.-A. (2004). Direction concepts in wayfinding assistance. In J. Baus, C. Kray, and R. Porzel (eds.), *Workshop on artificial intelligence in mobile systems (AIMS'04)* (pp. 1-8). Saarbrücken: SFB 378 Memo 84.

Klippel, A., Tappe, H., Kulik, L. and Lee, P.U. (submitted). Wayfinding choremes – A language for modelling conceptual route knowledge. (Submitted to Journal of Visual Languages and Computing).

Krieg-Brückner, B., Röfer, T., Carmesin, H.-O., and Müller, R. (1998). A taxonomy of spatial knowledge for navigation and its application to the Bremen autonomous wheelchair. In C. Freksa, C. Habel, and K.F. Wender (eds.), *Spatial Cognition. An interdisciplinary approach to representing and processing spatial knowledge* (pp. 373-397). Berlin: Springer.

Lovelace, K.L., Hegarty, M., and Montello, D.R. (1999). Elements of good route directions in familiar and unfamiliar environments. In C. Freksa & D Mark (eds.), *Spatial information theory. Cognitive and computational foundations of geographic information science* (pp. 65-82). Berlin: Springer.

Maaß, W. (1993). A cognitive model for the process of multimodal, incremental route descriptions. In A.U. Frank & I. Campari (Eds.), *Spatial information theory: A theoretical basis for GIS* (pp. 1-13). Berlin: Springer.

MacEachren, A.M. (1995). *How maps work*. New York: Guilford Press.

Miller, G.A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review* 63:81-97.

Montello, D.R. (in press). Navigation. In P. Shah & A. Miyake (eds.), *Handbook of visuospatial cognition*. Cambridge: Cambridge University Press.

Raubal, M. & Winter, S. (2002). Enriching wayfinding instructions with local landmarks. In M. Egenhofer & D. Mark (eds.), *Geographic information science - Second international conference GIScience 2002* (pp. 243-259). Berlin: Springer.

Richter, K.-F., Klippel, A., and Freksa, C. (2004). Shortest, fastest, — but what next? A different approach to route directions. In M. Raubal, A. Sliwinski, and W. Kuhn (eds.), *Geoinformation und Mobilität – von der Forschung zur praktischen Anwendung. Beiträge zu den Münsteraner GI-Tagen* (pp. 205-217). Münster: IfGI prints.

Schweizer, K., Katz, S., and Janzen, G. (2000). Orientierung im Raum - kognitive Grundlagen und sprachliche Realisierung. *Tourismus Journal* 4(1):79-101.

Shannon, C.E. & Weaver, W. (1949): *The mathematical theory of communication*. Urbana, Illinois: The University of Illinois Press.

Sorrows, M. & Hirtle, S.C. (1999). The nature of landmarks for real and electronic spaces. In C. Freksa & D. Mark (eds.), *Spatial information theory. Cognitive and computational foundations of geographic information science* (pp. 37-50). Berlin: Springer.

Tom, A. & Denis, M. (2003). Referring to landmarks or street information in route directions: What difference does it make? In W. Kuhn, M. Worboys, and S. Timpf (eds.), *COSIT 2003. Spatial information theory* (pp. 362-374). Berlin: Springer.

Tversky, B. & Lee, P. (1998). How space structures language. In C. Freksa, C. Habel, and K.F. Wender (eds.), *Spatial Cognition. An interdisciplinary approach to representing and processing spatial knowledge* (pp. 157-175). Berlin: Springer.

Tversky, B. & Lee, P. (1999). Pictorial and verbal tools for conveying routes. In C. Freksa & D.M. Mark (eds.), *Spatial information theory. Cognitive and computational foundations of geographic information science* (pp. 51-64). Berlin: Springer.

Wahlster, W., Baus, J., Kray, C., and Krüger, A. (2001). REAL: Ein ressourcenadaptives mobiles Navigationssystem. *Informatik Forschung und Entwicklung* 16:233-241.

Winter, S. (2003). Route adaptive selection of salient features. In W. Kuhn, M. Worboys, and S. Timpf (eds.), *COSIT 2003. Spatial information theory* (pp. 349-361). Berlin: Springer.