A Uniform Handling of Different Landmark Types in Route Directions

Kai-Florian Richter

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

Abstract. Landmarks are crucial for human wayfinding. Their integration in wayfinding assistance systems is essential for generating cognitively ergonomic route directions. I present an approach to automatically determining references to different types of landmarks. This approach exploits the circular order of a decision point's branches. It allows uniformly handling point landmarks as well as linear and areal landmarks; these may be functionally relevant for a single decision point or a sequence of decision points. The approach is simple, yet powerful and can handle different spatial situations. It is an integral part of GUARD, a process generating context-specific route directions that adapts wayfinding instructions to a route's properties and environmental characteristics. GUARD accounts for cognitive principles of good route directions; the resulting route directions reflect people's conceptualization of route information.

1 Introduction

Landmarks are crucial elements in human wayfinding. People use them to identify previously visited places and to orient themselves. Landmarks are a pertinent part of an environment's mental representation. They are also crucial in human direction giving. People refer to them frequently while providing another person with directions to reach a destination; they are mostly used to link actions to be performed to places in the environment. This holds for both verbal route directions and for sketch maps [1].

Integrating references to landmarks in wayfinding assistance systems is essential for automatically generating cognitively ergonomic route directions. We developed a generation process that implements cognitive principles of good route directions and allows for integrating different types of landmarks in instructions. To this end, features of an environment that may serve as a landmark need to be identified and it needs to be checked whether and how they are applicable in communicating the information necessary for route following. In this paper, I present an approach to testing the applicability of different types of landmarks in a uniform way; it is part of the generation process for route directions.

The paper is structured as follows: in the next section, I further illustrate the importance of landmarks for organizing spatial knowledge and their prominence in route directions. Section 3 introduces context-specific route directions, our approach to an automatic generation of route directions. In Section 4, I explain the prerequisites for automatically determining a landmark's functional role in route following; the determination

S. Winter et al. (Eds.): COSIT 2007, LNCS 4736, pp. 373–389, 2007.

is detailed in Section 5. Section 6 then demonstrates how context-specific route directions are generated and how different landmark types are integrated in the directions.

2 Landmarks in Route Directions

Many definitions of the term *landmark* can be found in the literature. Common to all, landmarks are defined as entities that are easily recognizable and memorizable. Presson and Montello provide a general definition stating that everything that stands out of the background may serve as a landmark [2]. In his seminal book 'The Image of the City', Lynch discusses why a feature may serve as a landmark: "Since the use of landmarks involves singling out of one element from a host of possibilities, the key physical characteristic of this class is singularity, some aspect that is unique or memorable in the context" ([3], pp. 78–79). Lynch further elaborates spatial prominence: elements may be established as landmarks either because they are visible from many locations or because of local contrast to nearby elements.

Sorrows and Hirtle pick up these considerations and list features that let a landmark stand out: *singularity*, *prominence*, *meaning*, and *prototypicality* [4]. Singularity applies to objects that are in sharp visual contrast with their surroundings. Prominence of entities refers to their spatial location in an environment; they are visible from many other locations or are located at a significant point, such as a major intersection. Some entities may be used as landmarks because they have a meaning common to many people stemming, for example, from their cultural or historical significance. Prototypicality is a characteristics similar to meaning in that such entities are referred to because they are typical representatives of a specific category. According to these characteristics, Sorrows and Hirtle identify three types of landmarks. A *visual landmark* is an entity used as landmark primarily because of contrast with its surrounding, because it has a prominent spatial location, or because its visual characteristics are easily memorizable. A *structural landmark* is defined as having a significant spatial role or location in the structure of space, while a *cognitive landmark* stands out because of its typical or atypical characteristics in the environment.

There is general agreement in the cognitive science literature that landmarks are eminently important for acquiring and organizing knowledge about our surrounding space. For example, there is evidence that spatial knowledge is organized hierarchically [5,6] around landmarks that function as anchor points for this knowledge [7,8]. In acquisition of space, landmarks are learned early on. The model of spatial knowledge acquisition of Siegel and White [9] sets landmarks to be the first kind of spatial knowledge acquired in a new environment. People recognize landmarks as places they have been before; connecting these places is the next step, termed *route learning*. Only in a final step, the different routes are integrated into survey knowledge that allows, for example, calculating shortcuts. Montello [10] questions this framework with regard to the (strict) order of these steps and states that some kind of metric knowledge, i.e. knowledge about distances, is acquired from the very beginning. But he does not question the important role landmarks play in spatial knowledge organization.

Landmarks are not only an important organizing concept for spatial knowledge, they also serve as navigational tool [11]. Landmarks identify decision points, origin and destination of a route, provide verification of route progress, provide orientation cues for homing vectors, and suggest regional differentiating features. This is echoed in research on route directions. People use landmarks to signal crucial actions, locate other landmarks in relation to the referenced landmark, and provide information to confirm that the right track is still followed [12,13,14]. The need to integrate landmarks into computational assistance systems has been confirmed by Tracy Ross and coworkers. In a set of usability studies they show that the use of landmarks in such systems significantly improves users' confidence in correctly executing the instructions and their navigation performance in both car navigation [15] as well as pedestrian wayfinding [16].

Landmarks may have different locations relative to a route. A landmark can either be at a decision point, at a route segment between two decision points [17], or in some distance to, but visible from the route (termed *distant landmark* [13]). Furthermore, there may be *global landmarks* [4]. These are outside the current environmental space, i.e. not immediately reachable by a wayfinder, but their location relative to the current space is known. Landmarks are usually assumed to be point entities. In Lynch's taxonomy of elements that structure spatial knowledge of a city, they are physical objects serving as point references that single out one element from a host of other objects [3]. This assumption is also prominent in research on the mental representation of spatial knowledge. In the model of Siegel and White [9], landmarks are strategic places that a wayfinder travels to and from to keep herself oriented. Hierarchies of spatial knowledge are formed based on point landmarks [5,6]. In their taxonomy of different types of landmarks, Sorrows and Hirtle [4] use the term landmark as it is defined by Lynch. Consequently, computational approaches that integrate landmarks also focus on point landmarks (e.g., [18,19]).

In our work on route directions, we extend the point notion of landmarks and integrate linear and areal entities, as well as structural elements, such as salient intersections [20,21,22]. Examples for linear landmarks include rivers or railway tracks; areal landmarks may be parks or big shopping malls. Typical salient intersections are T-intersections and roundabouts. With this extended view on landmarks, we comply with Presson and Montello's definition of a landmark as everything that sticks out from the background [2] (see also the anchor point theory [7]). Considering different types of landmarks leaves open more options to describe spatial situations, i.e. extends the expressibility of our representation. It also eases adapting instructions to the current situation. And it allows better accounting for human conceptualizations of spatial situations, i.e. it improves communication of the route directions.

3 Context-Specific Route Directions

We developed a computational process for generating route directions that takes different types of landmarks into account [23,20]. The process is termed GUARD, which stands for Generation of Unambiguous, Adapted Route Directions. This reflects that the process generates directions that unambiguously identify each route-segment and that adapt to the current action to be taken in the current surrounding environment. In that, we try to be as precise as possible with as little information as necessary; we take a Gricean perspective [24] in generating references. Our approach is in line with Dey's definition of context: "[...]any information that can be used to characterize the situation of an entity" ([25], p. 5). Accordingly, we coin the route directions generated by our process *context-specific route directions*.

For the adaptation, we need to account for the characteristics (the structure) of the environment in which route following takes place. The structure of an environment strongly influences the kind of instructions that can be given. The embedding of the route in the spatial structure surrounding it, the structure of that route itself, path annotations, and landmarks that are visible along the route all contribute to this influence. Based on an analysis of the basic constituents of routes and route directions, we identified classes of elements that can be used in route directions, as well as the spatial knowledge required to determine and interpret the corresponding references. We group these elements in a systematics on three different levels reflecting their relation to a specific route (see Table 1). They abstract to different degrees from a detailed description of a single decision point, which is captured by different levels of granularity. Global references abstract the most, while egocentric references and landmarks at decision points provide the most detailed description.

Table 1. The systematics of route direction elements. Elements are grouped on three levels according to their relation to the route: *Global References*, *Environmental Structure*, and *Path and Route*.

Global References	Environmental Structure	Path and Route
cardinal directions	edges	egocentric references
global landmarks	districts	landmarks at decision point
	slant	landmarks between decision points
		distant landmarks
		linear and areal landmarks
		path annotations

Based on the systematics, we implemented GUARD. It consists of four steps: in an initial step, every possible description of the action to be performed at each decision point is generated. These are represented in relational statements termed *abstract turn instructions* (ATI). The next step combines ATIs of consecutive decision points into a single instruction, performing so called *spatial chunking* [26]. As the initial chunking is simply syntax based, in the third step the generated chunks are checked against cognitive and representation-theoretic principles defining valid chunks. Finally, the fourth step generates the route directions in an optimization process. These steps are further illustrated in Section 6.

The inclusion of landmarks requires automatically determining their functional role in following a specific route. To this end, first, landmarks need to be assigned to those decision points they might be functionally relevant for. Second, it needs to be checked whether they unambiguously identify the route-segment to take at the decision points they are assigned to. The next section explains the basic representation used in the process and how landmarks are assigned to decision points; Section 5 illustrates how their applicability is checked.

4 Representing Routes and Landmarks

In our approach the environment's representation is coordinate-based. The street-network is represented as a graph, i.e. streets are represented as edges. Nodes that have a degree greater than three represent intersections. Nodes with a degree of two are used to reflect a street's geometry. Each node is associated with a coordinate. In the same coordinate system, features that may serve as landmarks are represented as points, lines, or polygons depending on their geometry.

A route is represented as a directed path in the graph. As already explained, of special interest are decision points, i.e. the nodes with a degree greater two. In the real world, these decision points are not just points. Here, a decision point denotes a certain area around an intersection which contains the point where the streets meet and part of the streets itself. It represents the configuration of streets at an intersection. That is, the branches meeting at a decision point are part of that point. This modeling corresponds to human conceptualization and reflects that humans usually turn in an extended process with deciding on the direction to turn and then changing their direction gradually while keeping their forward movement [1,27]. However, for assigning landmarks to decision points and for determining their applicability, we abstract from this gradual direction change. Decision points are modeled such that all branches meet in a single point (Fig. 1). This divides the area around an intersection into a *region before action* and a *region after action* [28].

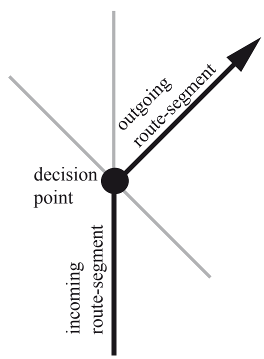

Fig. 1. A decision point with the two functionally relevant branches: the *incoming* and *outgoing* route-segment. The branches' meeting point is the actual decision point.

4.1 Assigning Landmarks to Decision Points

Having set the basic modeling of the environment, the landmarks represented in the annotated graph need to be functionally assigned to the route. As explained above, some landmarks may be functionally relevant for several decision points. Here, we need to distinguish two cases: the landmark spreads along the route which may, for example, hold for a linear landmark such as a river; or the landmark is distant to the route but visible from several decision points, for instance a church spire.

In the first case, the landmark's geometry is represented by a sequence of coordinate points. Even though belonging to the same landmark, these points can be individually assigned to a decision point. Different parts of such a landmark are relevant for different parts of the route, i.e. only that part of a landmark that is in the local surroundings of a decision point is functionally relevant for this decision point. Accordingly, for such landmarks, there is an exclusive assignment of coordinate points to decision points, i.e. each coordinate point is assigned to exactly one decision point. Landmarks distant to the route cannot be assigned exclusively to a specific decision point; this is independent of how they are geometrically represented. That is, next to categorizations of landmarks as, for example, discussed in [22] (see also Section 2), in the context of this work landmarks are distinguished between those that are only locally functionally relevant (i.e. that can be exclusively assigned to decision points) and those that are globally relevant. However, it is important to note that the role of a landmark may change along a route. It may serve as a distant landmark for several decision points, but be located at one of the route's decision points, i.e. serve as a local landmark there. In the following, I will concentrate on landmarks that are locally functionally relevant.

In order to assign these landmarks to decision points, a region around each decision point is defined. The size of this region depends on different parameters including travel mode and visibility[1]. Most importantly, the regions are chosen such that no other decision point falls within the region, i.e. the regions do not overlap. This way, we prevent conflicting assignment of a landmark's coordinate points to decision points; all coordinate points within the decision point's region are exclusively assigned to that decision point. This may lead to different sizes of the regions for different decision points. Decision points that are close together have smaller regions. This reflects that with decision points close together, a wayfinder has to decide on the further way to take more frequently, i.e. has to correctly interpret new spatial situations more often.

4.2 Ordering Information in Route Directions

For determining the functional role of a landmark, we exploit ordering information. Ordering information derives from the linear, planar, or spatial ordering of features [29]. It does not specify any further metric information, such as distances between these entities. This kind of information is a powerful structuring means and it is easy to determine as it only requires knowledge about a neighborhood relation between the relevant entities.

In routes, ordering is closely linked to orientation; a route, being a linear and directed (oriented) entity, induces an order on the entities along that route [30]. That is, the orientation of the route determines in which spatial and temporal order these entities are encountered. Furthermore, the configuration of entities at an intersection, for instance its branches or landmarks located there, can be described using circular ordering information. Figure 2 illustrates both these usages of ordering information.

[1] Checking for visibility is kept simple in the system. It is performed on the graph representation using scan-line methods. It just ensures that in the 2D projection an object is in line of sight from the route and in a distance shorter than some threshold.

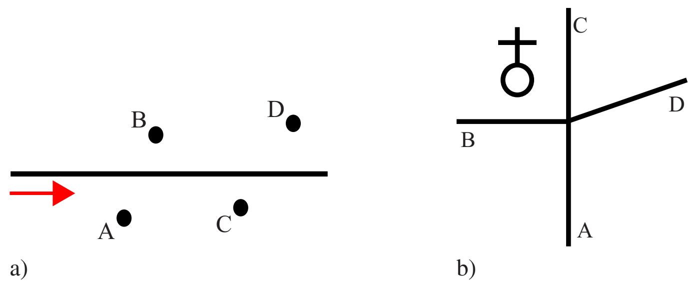

Fig. 2. a) Linear order of entities induced by the orientation of a route. The route is directed from left to right; the order is $A < B < C < D$. b) Circular order of entities at an intersection: starting with A, the order is $A < D < C < Church < B < A$.

5 A Uniform Handling of Landmarks in Route Directions

In the following, I will detail a uniform approach dealing with different types of landmarks, As a basic case, I start with point landmarks at a decision point. This case illustrates the general principles; ongoing from there, extensions needed for handling other types of landmarks will be explained. Generally, the location of a landmark relative to a decision point is required to appropriately reference the landmark in route directions; the relation to incoming and outgoing route-segment and, consequently, to their meeting point (the decision point itself) needs to be determined.

5.1 Point Landmarks at Decision Points

For point landmarks at a decision point, there are three locations that can be distinguished functionally. The turning action may either occur *after* passing a landmark ("turn left after the church"), it may take place *before* the landmark is passed ("turn left before the church"), or the landmark may be *at* the decision point, but not located at a functionally relevant branch ("turn left at the church") (see Fig. 3). In other words, the landmark may be next to the incoming route-segment, next to the outgoing route-segment, or next to any of the other branches of the decision point. In order to

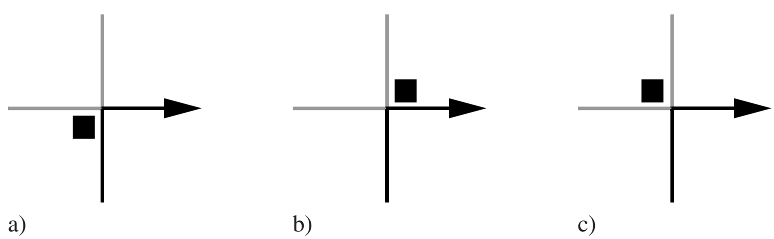

Fig. 3. Three functionally different locations of a landmark at a decision point: a) turning action after passing the landmark; b) turning action before passing the landmark; c) landmark not at a functionally relevant branch

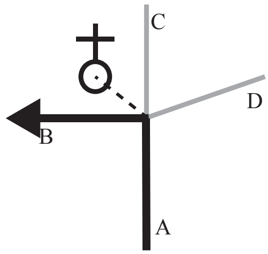

Fig. 4. A virtual branch integrating a landmark at a decision point into the circular order of the decision point's branches

determine the location of the landmark relative to the decision point, we need to determine this *next to* relation. It corresponds to a neighborhood relation of landmark and branch. This neighborhood can be extracted from the circular ordering of the decision point's branches. We introduce a virtual branch ranging from the decision point to the point landmark. This way, the landmark becomes part of the branches' circular ordering (see Fig. 4).

We can now determine whether the virtual branch is neighbored to one of the functionally relevant branches. Two branches are neighbored if one succeeds the other in the ordering. If the virtual branch representing the landmark's location is direct successor or predecessor of the incoming route-segment, the turning action is performed after the landmark is passed. We represent this situation with the relation $\texttt{lm}^{<}$. The turning action is performed before the landmark if in the ordering the virtual branch is neighbor to the outgoing route-segment, represented by $\texttt{lm}^{>}$. All situations in which the landmark is not neighbored to a functionally relevant branch are represented by $\texttt{lm}^{-}$. To ensure that the landmark is really located before the decision point, two additional virtual branches perpendicular to the incoming route-segment are introduced. These branches demarcate the before-region. Without them, a landmark may be neighbored to the incoming route-segment though it is located after the decision point (see Fig. 5). The same holds for the after-region correspondingly. For relation $\texttt{lm}^{>}$, further restrictions apply with respect to possible ambiguities in identifying a branch. This is reported in [28], along with further details on this basic case—point landmarks at a decision point.

5.2 Other Landmark Types at Decision Points

Next, we consider features of the environment whose landmark character is dominated by a specific part that may not be visible from everywhere around the feature. This holds, for example, for buildings housing shops that have a salient store-front [18]. These features may be represented as points, as well. But, additionally, the salient side (the façade) needs to be captured; this is done by adding a vector pointing in direction of the salient side. The relation of the landmark to the decision point is determined as illustrated above. Then, the intersection of the vector representing the façade with the functionally relevant branches is calculated. If the landmark is next to the incoming route-segment, the vector needs to intersect with this route-segment for the landmark

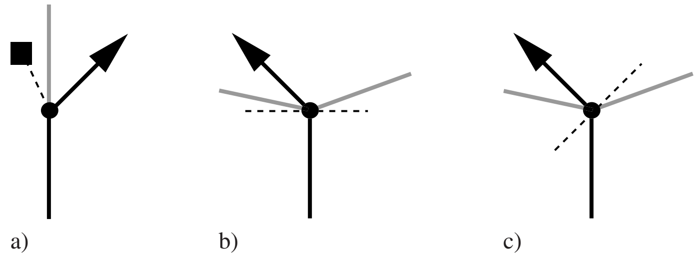

Fig. 5. a) A spatial situation where the landmark is located next to the incoming route-segment, but after the decision point; $\mathtt{lm}^{<}$ is not a valid result here. b) Virtual branches demarcating the before-region… c) …and the after-region.

to be applicable. The same holds for the outgoing route-segment; the vector needs to intersect it to reference the landmark at the decision point.

The basic principle can also be used for landmarks that are not represented as point features. The three relations $\mathtt{lm}^{<}$, $\mathtt{lm}^{>}$, and $\mathtt{lm}^{-}$ are still sufficient to handle these cases. A landmark's geometry is defined by a sequence of coordinates. For each coordinate point it is possible to determine its relation to a decision point by introducing a virtual branch connecting it to the decision point and then calculating the relation as described above. The functional role of a linear or areal landmark in route following is then determined by the resulting sequence of relations holding for its coordinate points.

For example, the location of a landmark at a decision point that cannot be represented as a single point (see Fig. 6) is addressable in the same way as a point landmark, i.e. there are three functionally different locations (turning action before the landmark, after the landmark, or at the landmark). If the resulting relation is $\mathtt{lm}^{<}$ for all coordinate points, the landmark is passed before the turning action occurs. We can assign $\mathtt{lm}^{<}$ to denote the location of the landmark with respect to the decision point. If all relations are $\mathtt{lm}^{>}$, it is passed after the turning action, and the relation for the landmark is $\mathtt{lm}^{>}$. If the resulting relations differ (especially if some are $\mathtt{lm}^{-}$) the landmark is not completely located in before- or after-region. Therefore, $\mathtt{lm}^{-}$ is the adequate relation to represent the landmark's location.

5.3 Extended Landmarks Relevant for Several Decision Points

Linear and areal landmarks may stretch along the route, i.e. be applicable for several consecutive decision points. As detailed in Section 4.1, the coordinate points defining their geometry get assigned to different decision points. This way, it is possible to determine for individual decision points the relation to the relevant part of a linear or areal landmark. Using the same approach as before, we can get from the relation of single coordinate points to a decision point to the relation of a landmark to a single decision point. And from there, we can get to the relation of that landmark to a sequence of decision points.

It may, for example, be possible to refer to a river to indicate the direction to take for several decision points in instructions as "follow the river until the gas station". In this case, the river must unambiguously indicate the direction to take at each decision

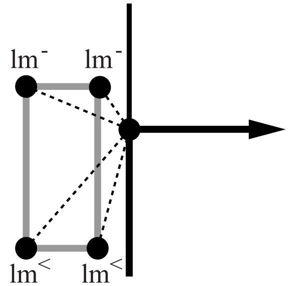

Fig. 6. Example of an extended landmark at a decision point. For each coordinate point, its relation to the decision point is determined. The resulting relation of the landmark to the decision point is $\mathtt{lm}^-$ as different relations hold for the points.

point up to the one where the gas station is located. This is the case if for each of these decision points the river is both next to the incoming and the outgoing route-segment. To check for this, for all coordinate points that are assigned to the current decision point the relation to the decision point is determined. This is done in movement direction, i.e. the coordinate points are processed in the order they would be passed along the route. For the linear landmark to be applicable, for the first l coordinate points relation $\mathtt{lm}^<$ needs to hold, for the last m coordinate points $\mathtt{lm}^>$ needs to hold ($l, m \geq 1$). No coordinate point may be in relation $\mathtt{lm}^-$, and there may be no switch back to $\mathtt{lm}^<$ after $\mathtt{lm}^>$ has been assigned to a coordinate point. If this condition is fulfilled the linear landmark can be used to indicate the direction to take at the decision point. We denote this with the relation `follow`. With areal landmarks, the same method can be applied to check whether it may be used to guide a wayfinder around that landmark along the route. As long as one part of the coordinate points is next to the incoming route-segment and the other part next to the outgoing route-segment, the route passes along the areal landmark.

In conclusion, the basic principle of determining the location of a point landmark relative to a decision point and the distinction of the three different relations $\mathtt{lm}^<$, $\mathtt{lm}^>$, and $\mathtt{lm}^-$ is sufficient to handle different spatial situations involving different types of landmarks (namely point, linear, and areal landmarks). The situations covered so far all involve passing a landmark, either at a single decision point or walking along it for several consecutive decision points. A simple extension of the basic formalism allows for handling further spatial situations in which the involved landmark's role may be more complex. When crossing a river, for example, the landmark is both to the left and the right of the route. That is, in the corresponding representation of the route, part of the coordinate points representing the landmark's geometry are on the one side of the route-segment, the other part on the other side.

This sideness can be captured using the ordering of branches, as well. The branches' ordering (including the virtual branch connecting a coordinate point) always gets calculated in counterclockwise direction. That is, the order of the branches is ordered, as well. Taking this information into account then allows inferring the sideness of a coordinate point since we know the movement direction along a route-segment; for the incoming

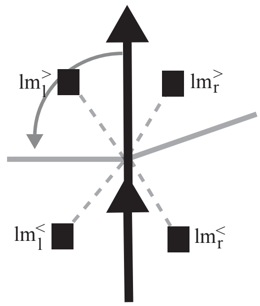

Fig. 7. Distinguishing the sideness of a landmark with respect to a functionally relevant route-segment. The gray arrow shows the direction of determining the branches' order.

route-segment the predecessor in the ordering is left of the branch, the successor right of the branch. For the outgoing route-segment, it is vice versa; the predecessor is right of the branch, the successor left of it (see Fig. 7). We denote this with the indices l and r and end up with five different relations: $\mathtt{lm}_\mathtt{l}^<$, $\mathtt{lm}_\mathtt{r}^<$, $\mathtt{lm}_\mathtt{l}^>$, $\mathtt{lm}_\mathtt{r}^>$, $\mathtt{lm}^-$.

Using these relations, we can capture additional spatial situations. Picking up the example of crossing a linear landmark, a possible sequence of relations that corresponds to this situation is the following: there are l coordinate points for which the relation $\mathtt{lm}_\mathtt{l}^<$ holds, followed by m coordinate points with relation $\mathtt{lm}_\mathtt{r}^<$. More generally, crossing a linear landmark is represented by an arbitrary number of coordinate points in relation $\mathtt{lm}^-$, followed by l coordinate points in exactly one relation of the set $\mathtt{lm}_\mathtt{l}^<, \mathtt{lm}_\mathtt{r}^<, \mathtt{lm}_\mathtt{l}^>, \mathtt{lm}_\mathtt{r}^>$, then m coordinate points being next to the same route-segment but on the other side, for example, switching from $\mathtt{lm}_\mathtt{l}^<$ to $\mathtt{lm}_\mathtt{r}^<$, followed again by an arbitrary number of coordinate points in relation $\mathtt{lm}^-$ (with $l, m \geq 1$).

These examples demonstrate that it is possible to define patterns of relations that capture the functional role of different landmark types in route following. This is based on a simple principle of using ordering information to determine the location of a point landmark at a decision point and by distinguishing five relational terms to capture the functionally relevant possible locations. The presented approach allows identifying spatial situations and generating appropriate references to landmarks. It is important to note that this is not meant to be a model of how humans judge spatial situations and of how they determine adequate descriptions for these situations.

6 Generating Context-Specific Route Directions

The method for determining the applicability of different landmark types is an integral part of GUARD, the process for generating context-specific route directions. As stated in Section 3, this process automatically generates route directions that exploit different features of an environment; it is realized in four steps. In this section, I will provide more details on each step and will give an example for route directions generated by the model. Figure 8 depicts an example route that has been put together using different real-world intersections whose data has been collected as input to GUARD. The following illustration is based on this route.

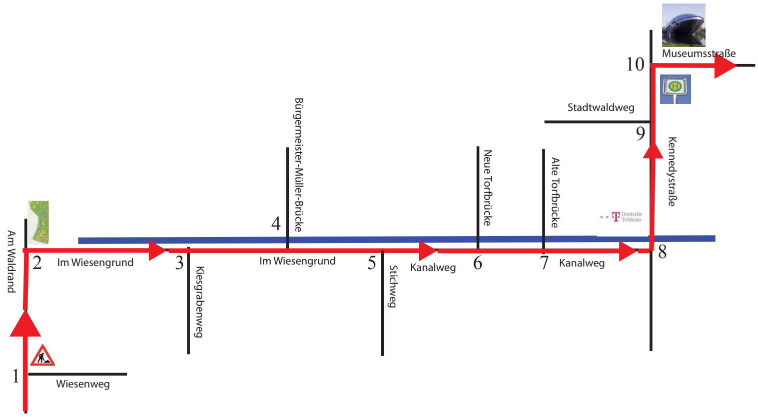

Fig. 8. An example route; the decision points are numbered for future reference

6.1 Generating Abstract Turn Instructions

In the first step of the generation process, for every decision point of the route all possible instructions are generated. To this end, for every element listed in the systematics (see Table 1), it is checked whether it is applicable at all and whether it results in an instruction that unambiguously describes the further direction to take. In the abstract turn instructions, actions to be performed at a decision point are represented as *direction relations*; each kind of action is denoted by a specific relational term. Actions result in a change of heading (which might be 0 if continuing straight on). In the graph representation, this change can be measured as an angular deviation from the previous heading.

For egocentric references, this angle is matched to a category representing a turning direction (e.g., *left*) to determine which relation to use. These categories are represented as angle intervals; we use the direction model of [31] that distinguishes seven different directions. For the first decision point of the example route, for instance, the angular deviation from the current movement direction is in an interval representing going straight, denoted by the relation `straight`. There is also a construction site, modeled as point landmark. Using the mechanism presented in Section 5.1 results in $\mathtt{lm}^{>}$, i.e. the landmark is located after the turning action. Possible abstract turn instructions for the first decision point are $(\mathtt{DP}_1,\ \mathtt{straight})$ and $(\mathtt{DP}_1,\ \mathtt{straight\ construction\text{-}site\ lm}^{>})$.

Other kinds of elements, for instance linear landmarks, do not rely on the angular deviation in providing information on the further way to take. The river next to considerable part of the route is a good example. The coordinate points representing the river get exclusively associated with decision points three to eight, respectively. For each of these decision points, one part of the assigned coordinate points are in the before-region and the other in the after-region; checking the landmark's applicability results in the relation `follow`. Accordingly, the landmark is suited to indicate the further direction to take at these decision points. All in all, the first step of the generation process results in a set of possible instructions for each decision point.

6.2 Spatial Chunking

Spatial chunking is a process that combines the actions to be performed at several consecutive decision points into a single action [26,19]. This is a crucial process in conceptualizing routes, i.e. in forming a mental representation of a route. Humans also commonly apply spatial chunking in giving route directions.

GUARD performs spatial chunking in two steps (which cover steps two and three of the generation process). In an initial step, abstract turn instructions are combined based on two simple syntactic rules: first, ATIs that employ the same direction relation can be combined; second, the egocentric direction relation `straight` can be combined with any other egocentric direction relation. The first rule corresponds to a "do n-times"– or "do until"–instruction, for instance "turn three times right"; the second rule covers situations where going straight at a decision point does not need to be explicitly stated, for instance, in "turn left at the third intersection". Both rules ensure that references to landmarks are sensibly handled, i.e. that each chunk only refers to a single landmark.

Then, it is checked whether the chunks generated using these rules are cognitively and structurally sensible. For example, instructions that refer to following linear landmarks require information on the point along the route upon which the landmark is to be followed. This information, termed *end qualifier*, typically is provided by a point landmark at a decision point. Thus, one of the chunking principles used in this step states that chunks based on linear landmarks need to end with a decision point where another landmark is present. Generally, in this step chunks are pruned until they adhere to the defined cognitive principles.

Performing spatial chunking as a two-step process is advantageous as it allows for more flexibility; different cognitive chunking principles can be implemented and used with the model while always using the same basic chunking mechanism.

6.3 Optimization

The actual generation of the route directions is performed in the fourth step. Here, from all possible abstract turn instructions for each decision point the one that is best gets selected. This is an optimization process; accordingly, 'best' is defined relative to an optimization criterion. Just as with the chunking principles, the model is flexible with respect to the optimization criterion used. In the literature, different principles for good route directions can be found that can be turned into optimization criteria (e.g., [12,13]). A straightforward criterion is to aim for the minimum number of chunks. This reduces the cognitive load of a wayfinder as the communicated route directions get shorter and, hence, the wayfinder needs to remember less information. And the wayfinder's processing of the route directions is also reduced since the instructions are already chunked.

There are two optimization approaches implemented in the model. Local optimization proceeds through the route from origin to destination. For each decision point, it selects the best chunk starting with this decision point. It then checks whether this chunk has to be integrated into the current partial route directions, i.e. whether this chunk improves the directions. If this is the case, the directions are rebuild using this chunk. Local optimization does not necessarily find an optimal solution. Global optimization guarantees this by generating every possible solution (every possible combination of chunks). This is computationally more complex than local optimization. However, computational tests show that in practice both approaches almost always result in the equally optimal route directions.

Table 2 shows the resulting context-specific route directions for the example route of Figure 8 as an XML-specification. Each chunk is grouped between <Direction>-tags. For each chunk, a verbal externalization is listed.[2]

Table 2. The resulting context-specific route directions and their verbal externalization. On the left, the part of the route that is covered by each chunk is depicted (rotated to fit in the table).

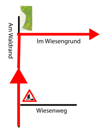	<Direction> <instruction point="1" relation="straight"> </instruction> <instruction point="2" relation="right"> landmark id="1" type="Map" relation="lm<"/> </instruction> </Direction>	"At the next junction go straight, and then turn right before a map post."
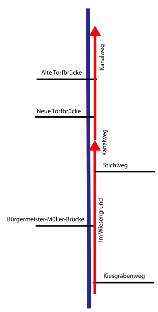	<Direction> <instruction point="3" relation="follow"> <landmark id="100" type="River"/> </instruction> <instruction point="4" relation="follow"> <landmark id="100" type="River"/> </instruction> <instruction point="5" relation="follow"> <landmark id="100" type="River"/> </instruction> <instruction point="6" relation="follow"> <landmark id="100" type="River"/> </instruction> <instruction point="7" relation="follow"> <landmark id="100" type="River"/> <endqualifier id="7" type="Signage" name="Telekom Schild"/> </instruction> <instruction point="8" relation="left"> <landmark id="7" type="Signage" name="Telekom Schild" relation="lm<"/> </instruction> </Direction>	"Keep following the river until a Telekom sign and then, turn left after the Telekom sign."
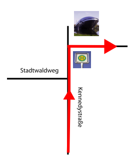	<Direction> <instruction point="9" relation="towards"> <landmark id="9" type="SightSeeingObject" name="Universum"/> <endqualifier id="12" type="BusStop"/> </instruction> <instruction point="10" relation="right"> <landmark id="12" type="BusStop" relation="lm<"/> </instruction> </Direction>	"Go towards the Universum until a bus stop, and then, turn right after the bus stop."

To test the performance of context-specific route directions, we performed an initial user-test as a desktop study. Subjects had to, first, rate context-specific route directions and, second, reach a destination location after memorizing route directions by clicking on the branch to take on intersection photographs. Subjects rated context-specific route directions as useful and well understandable, and even though the setup of the study is somewhat artificial, subjects were well able to reach the destination using context-specific route directions.

[2] Verbal externalization of context-specific route directions is not part of the process directly; it is done in collaboration with Prof. John Bateman, Universität Bremen, using their natural language generation system.

7 Conclusions

In this paper I presented an approach to automatically determining the functional role of landmarks in route following. The approach allows integrating different types of landmarks in automatically generated route directions. The circular order that the branches of a decision point form and the order of events in route following that is induced by the directedness of a route are exploited. Distinguishing five different relations, the location of a landmark's coordinate points relative to a decision point are determined. The sequence of these relative locations are compared to patterns representing possible functional roles of the landmark. This way, it is possible to cover different spatial situations based on a single simple principle. The approach is an integral part of GUARD, the generation process for context-specific route directions that automatically generates route directions that adapt to a route's properties and environmental characteristics. Generation of the route directions is a four-step process; optimization is used to decide on which description to choose.

In this process, we integrate different types of landmarks. This is a crucial step in automatically generating cognitively ergonomic route directions. The process also performs spatial chunking which is another basic principle of human direction giving. However, currently it depends on landmarks being known. Except for structural landmarks (e.g., T-intersections), potential landmarks are not identified from a given environmental representation. Integrating approaches that aim for automatically identifying landmarks (e.g., [18,32]) would be another important step for a truly automatic generation of route directions. Another aspect of future work is using the model to find the optimal route applying route selection criteria related to both the difficulty of describing the route and the difficulty of navigating the route (in line with, e.g., [33]), instead of finding the optimal description for a given route as it is done by now. Furthermore, due to its flexibility the model may serve as test-bed for empirical studies testing different principles of generating route directions.

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

References

1. Tversky, B., Lee, P.U.: Pictorial and verbal tools for conveying routes. In: Freksa, C., Mark, D.M. (eds.) Spatial Information Theory - Cognitive and Computational Foundations of Geopraphic Information Science, pp. 51–64. Springer, Berlin (1999)
2. Presson, C.C., Montello, D.R.: Points of reference in spatial cognition: Stalking the elusive landmark. British Journal of Developmental Psychology 6, 378–381 (1988)
3. Lynch, K.: The Image of the City. MIT Press, Cambridge (1960)
4. Sorrows, M.E., Hirtle, S.C.: The nature of landmarks for real and electronic spaces. In: Freksa, C., Mark, D.M. (eds.) Spatial Information Theory - Cognitive and Computational Foundations of Geopraphic Information Science, pp. 37–50. Springer, Berlin (1999)
5. Sadalla, E.K., Burroughs, J., Staplin, L.J.: Reference points in spatial cognition. Journal of Experimental Psychology: Human Learning and Memory 6(5), 516–528 (1980)

6. Hirtle, S.C., Jonides, J.: Evidence of hierarchies in cognitive maps. Memory & Cognition 13(3), 208–217 (1985)
7. Couclelis, H., Golledge, R.G., Gale, N., Tobler, W.: Exploring the anchor-point hypothesis of spatial cognition. Journal of Environmental Psychology 7, 99–122 (1987)
8. McNamara, T.P.: Spatial representation. Geoforum 23(2), 139–150 (1992)
9. Siegel, A.W., White, S.H.: The development of spatial representations of large-scale environments. In: Reese, H. (ed.) Advances in Child Development and Behaviour, pp. 9–55. Academic Press, New York (1975)
10. Montello, D.R.: A new framework for understanding the acquistion of spatial knowledge in large-scale environments. In: Egenhofer, M.J., Golledge, R.G. (eds.) Spatial and Temporal Reasoning in Geographic Information Systems, pp. 143–154. Oxford University Press, New York (1998)
11. Golledge, R.G.: Human wayfinding and cognitive maps. In: Golledge, R.G. (ed.) Wayfinding Behavior — Cognitive Mapping and Other Spatial Processes, pp. 5–46. John Hopkins University Press, Baltimore, USA (1999)
12. Denis, M.: The description of routes: A cognitive approach to the production of spatial discourse. Cahiers Psychologie Cognitive 16(4), 409–458 (1997)
13. Lovelace, K.L., Hegarty, M., Montello, D.R.: 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 Geopraphic Information Science, pp. 65–82. Springer, Berlin (1999)
14. Michon, P.E., Denis, M.: When and why are visual landmarks used in giving directions? In: Montello, D.R. (ed.) Spatial Information Theory - Cognitive and Computational Foundations of Geopraphic Information Science, pp. 400–414. Springer, Berlin (2001)
15. May, A.J., Ross, T., Bayer, S.H., Burnett, G.: Using landmarks to enhance navigation systems: Driver requirements and industrial constraints. In: Proceedings of the 8th World Congress on Intelligent Transport Systems, Sydney, Australia (2001)
16. Ross, T., May, A., Thompson, S.: The use of landmarks in pedestrian navigation instructions and the effects of context. In: Brewster, S., Dunlop, M. (eds.) Mobile Computer Interaction – MobileHCI 2004, pp. 300–304. Springer, Berlin (2004)
17. Herrmann, T., Schweizer, K., Janzen, G., Katz, S.: Routen- und Überblickswissen — konzeptuelle Überlegungen. Kognitionswissenschaft 7(4), 145–159 (1998)
18. Raubal, M., Winter, S.: Enriching wayfinding instructions with local landmarks. In: Egenhofer, M.J., Mark, D.M. (eds.) GIScience 2002. LNCS, vol. 2478, pp. 243–259. Springer, Heidelberg (2002)
19. Dale, R., Geldof, S., Prost, J.P.: Using natural language generation in automatic route description. Journal of Research and Practice in Information Technology 37(1), 89–105 (2005)
20. Richter, K.F., Klippel, A.: A model for context-specific route directions. In: Freksa, C., Knauff, M., Krieg-Brückner, B., Nebel, B., Barkowsky, T. (eds.) Spatial Cognition IV, pp. 58–78. Springer, Berlin (2005)
21. Klippel, A., Richter, K.F., Hansen, S.: Structural salience as a landmark. In: Workshop Mobile Maps 2005, Salzburg, Austria (2005)
22. Hansen, S., Richter, K.F., Klippel, A.: Landmarks in OpenLS - a data structure for cognitive ergonomic route directions. In: Raubal, M., Miller, H.J., Frank, A.U., Goodchild, M.F. (eds.) GIScience 2006. LNCS, vol. 4197, pp. 128–144. Springer, Heidelberg (2006)
23. Richter, K.F., Klippel, A., Freksa, C.: Shortest, fastest, - but what next? a different approach to route directions. In: Raubal, M., Sliwinski, A., Kuhn, W., (eds.) Geoinformation und Mobilität - von der Forschung zur praktischen Anwendung. Beiträge zu den Münsteraner GI-Tagen 2004 IfGIprints, Institut für Geoinformatik, Münster pp. 205–217 (2004)
24. Grice, H.P.: Logic and conversation. In: Cole, P., Morgan, J.L. (eds.) Speech Acts. Syntax and Semantics, vol. 3, pp. 41–58. Academic Press, New York (1975)

25. Dey, A.K.: Understanding and using context. Personal and Ubiquitous Computing 5(1), 4–7 (2001)
26. Klippel, A., Tappe, H., Habel, C.: Pictorial representations of routes: Chunking route segments during comprehension. In: Freksa, C., Brauer, W., Habel, C., Wender, K.F. (eds.) Spatial Cognition III, pp. 11–33. Springer, Berlin (2003)
27. Klippel, A.: Wayfinding Choremes — Conceptualizing Wayfinding and Route Direction Elements. PhD thesis, Universität Bremen (2003)
28. Richter, K.F., Klippel, A.: Before or after: Prepositions in spatially constrained systems. In: Proceedings of Spatial Cognition, Berlin, Springer (to appear)
29. Schlieder, C.: Reasoning about ordering. In: Frank, A.U., Kuhn, W. (eds.) Spatial Information Theory - A Theoretical Basis for GIS, Springer, Berlin (1995)
30. Pederson, E.: How many reference frames? In: Freksa, C., Brauer, W., Habel, C., Wender, K.F. (eds.) Spatial Cognition III, pp. 287–304. Springer, Berlin (2003)
31. Klippel, A., Dewey, C., Knauff, M., Richter, K.F., Montello, D.R., Freksa, C., Loeliger, E.A.: Direction concepts in wayfinding assistance systems. In: Baus, J., Kray, C., Porzel, R., (eds.) Workshop on Artificial Intelligence in Mobile Systems (AIMS'04), SFB 378 Memo 84; Saarbrücken, pp. 1–8 (2004)
32. Elias, B.: Extracting landmarks with data mining methods. In: Kuhn, W., Worboys, M., Timpf, S. (eds.) Spatial Information Theory: Foundations of Geographic Information Science, pp. 375–389. Springer, Berlin (2003)
33. Duckham, M., Kulik, L.: Simplest paths: Automated route selection for navigation. In: Kuhn, W., Worboys, M., Timpf, S. (eds.) Spatial Information Theory: Foundations of Geographic Information Science, pp. 169–185 (2003)