Landmark Hierarchies in Context

Stephan Winter
Department of Geomatics
The University of Melbourne, Australia
winter@unimelb.edu.au

Martin Tomko[1]
CRCSI, Department of Geomatics
The University of Melbourne, Australia
m.tomko@pgrad.unimelb.edu.au

Birgit Elias[2]
Institute of Cartography and Geoinformatics
University of Hannover, Germany
birgit.elias@ikg.uni-hannover.de

Monika Sester
Institute of Cartography and Geoinformatics
University of Hannover, Germany
monika.sester@ikg.uni-hannover.de

[1]Supported by the Cooperative Research Centre for Spatial Information, whose activities are funded by the Australian Commonwealth's Cooperative Research Centres Programme.
[2]Supported by the German Research Foundation (DFG).

Abstract

We are interested in the generation of distinguishing place or route descriptions for urban environments. Such descriptions require a hierarchical model of the discourse, the elements of the city. We postulate that cognitive hierarchies, as used in human communication, can be sufficiently reflected in machine-generated hierarchies. In this paper we (a) propose a computational model for the generation of a hierarchy of one of these elements of the city, landmarks, and (b) demonstrate that a set of filter rules applied on this hierarchy derives distinguishing route descriptions from spatial context.

Keywords: navigation, landmarks, hierarchies, route directions, place descriptions.

1 Introduction

When communicating routes or place descriptions people refer typically to landmarks. These landmarks are cognitive anchors of the learned environment that are assumed to be shared with other people, or at least with participants of a particular conversation. Furthermore, there is evidence for a hierarchical structure of cognitive spatial representations [36, 13], from which one can infer a hierarchical system of landmarks, distinguishing, let us say, landmarks representing the city from landmarks representing a street intersection. This hierarchy is also reflected in route or place descriptions, if the speaker assumes that the listener shares some knowledge of the environment [33, 27, 2].

If we want to design geographic information services, such as navigation services or location-based services, that are able to communicate about space hierarchically, we need to answer the following questions:

1. What are the driving forces for a hierarchical organization of landmarks, and how can we automatically organize a set of landmarks in a hierarchical system?
2. What are the specific properties of a ranked order? Is it a levelled hierarchy? Can we, for example, find ontological or at least operational specifications of levels, and is the mapping of a landmark into this system unique?
3. What about the duality of landmarks and their reference regions [3, 23] – is there a dual hierarchy of regions? What are the properties of this dual hierarchy, and if the mapping of landmarks is not unique, i.e., if individual landmarks can occur at different levels in the hierarchy, do their reference regions vary with the level? In other words, what is the spatial context of a landmark?
4. Finally, given a hierarchy of landmarks, how does a geographic information service make use of it? What are the rules for transiting between the levels in a communication process? How are they reflected in the dual hierarchy?

In this paper, we try to answer these questions. The underlying hypothesis is that we can find a formal model for a machine-generated hierarchy of landmarks from geographic datasets. Ideally the ranking criteria in the model are motivated by cognitive relevance.

We start with collecting previous work on landmark identification and the use of landmarks in route descriptions (Section 2). The known methods of identifying landmarks automatically from datasets all apply regionalized approaches, which makes sense on the basis that an object can only become a landmark relative to other objects in that region. We take advantage of the regionalization, and exploit it to derive hierarchies of landmarks from examples. The observations help us to derive some general properties of a model for such hierarchies (Section 3). In Section 4 we present a formal model and algorithm to determine landmark hierarchies. Then we propose a procedure to identify and order landmarks for a limited domain, buildings. The procedure will be applied to a test area to demonstrate its behavior (Section 5). Finally we adapt an existing formal model for generating route or place descriptions from region hierarchies to landmark hierarchies, with respect to spatial context (Section 6). We expect to find a complementary hierarchy to the hierarchies of other elements of the city [21] that can be used, either on its own or in combination, for hierarchic route or place descriptions. The paper concludes with a discussion and an outlook in Section 7.

2 Previous work

In this section we review the relevant literature on landmarks, hierarchic organization of cognitive space, and spatial hierarchies.

2.1 Landmarks

Landmarks are features in the physical environment that are relatively better known and define the location of other points [28]. Since landmarks can be experienced, they are represented in cognitive representations of space as anchor points [3]. Because of their special role in cognitive representations—they may even be the first class citizens of cognitive representations—they are particularly suited for human-to-human communication [6, 14].

Landmarks are one of the five elements Lynch [21] has identified in the image of a city. These are *place*, *path*, *barrier*, *district*, and *landmark*. Today it is a common view that any element of a city can have landmark quality according to the above definition [e.g., 31]. However, in this paper we limit ourselves to Lynch's more narrow view, and consider landmarks as point-like elements in the city.

Literature provides ample evidence that cognitive representations of space are hierarchical [36, 13]. Hierarchic representations turn out to be efficient in cognitive spatial reasoning and communication. For example, place descriptions based on hierarchies are shorter, and hence, more relevant compared to descriptions on one level of granularity [33, 27, 26]. Hence, computational linguistics proposed hierarchical referring expressions [2, 5]. Route descriptions can be generated using the same principles [45], and we will come back to this point in Section 6. Another example for the efficiency of hierarchic spatial reasoning is the route generation itself [46, 42], which motivated the development of hierarchical spatial data models for use in artificial intelligence [40, 17]. In the cartographic context, the metaphor of a *cognitive zoom* was developed on these data models [10, 37].

2.2 Categorization of landmarks

Landmarks are frequently categorized into global and local landmarks. Since landmarks are used in two distinct contexts—wayfinding and representation of spatial knowledge [28]—the meaning of this categorization varies.

In a wayfinding context, global landmarks are typically used for conveying directional information; they are at a distance, or off the route. Local landmarks are typically used for conveying positional information; they are close to the route, and are sometimes further categorized into landmarks at decision points, landmarks at potential decision points, and on-route landmarks along segments [20, 15].

In the context of spatial knowledge, landmarks are considered as anchors of cognitive representations of space. In this context, a local landmark is a landmark that refers to a region defined by vista space, and a global landmark is a landmark that refers to a larger region [22]. With this weak definition at hand, we can already assume that there will be several levels of global landmarks in a city, distinguishable only by the size of their reference region. A sharper definition will be possible after formalizing the reference regions.

The notions of landmark hierarchies and global landmarks in spatial knowledge are related to an observation in game theory [32]: if two persons have to meet in a reference region (i.e., in a given spatial context) but cannot communicate with each other to arrange a meeting, these persons will likely decide to go to its landmark. Additionally they choose a landmark in the temporal domain for the expected meeting time. This most salient point in space-time is nowadays called *Schelling point*.

3 Landmark hierarchies

In this section we show with examples that (point-like) landmarks in a city have a cognitive hierarchy, forming a partial order. We also show that the landmarks have a dual in their reference regions. Consequently the reference regions form the same hierarchy.

3.1 Motivational example

People naturally rank landmarks. For example, when entering a taxi at Melbourne's airport Tullamarine, it makes perfect sense to tell the taxi driver "To the Royal Hospital, then turn right, and you will see a 7-Eleven" (Figure 1). The Royal Hospital, located on the edge of Royal Park, is in a suburb next to the central business district, and about 20km from Tullamarine. Being the premier health care provider in Victoria, the hospital is well-known, and one can assume that the taxi driver will know it – and find the way. Note that other cities in Australia have a Royal Hospital as well, but this seems not to concern the taxi driver in this situation. However, Melbourne has several hospitals, and hence the passenger felt impelled to name the specific one [45]. In contrast, the 7-Eleven would not have worked as the first reference. The company logo and façade makes it distinct in its neighborhood, but business directories list 259 7-Eleven shops in Melbourne. Hence, a 7-Eleven is a local landmark, distinguished only in vista environments, and in the given expression it is dependent on the hospital.

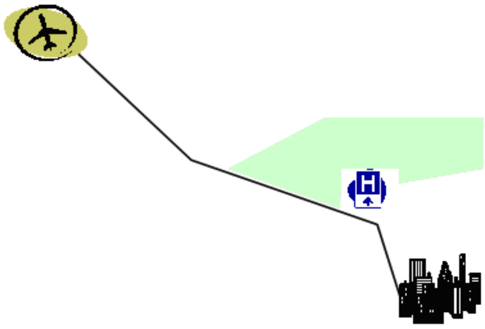

Figure 1: Melbourne's Royal Hospital, on the edge of Royal Park, is close to the Central Business District, and about 20km from Melbourne's international airport, Tullamarine.

Now consider a phone conversation I have with my friend. She is currently in the Royal Park and I want to meet her in a café close to the Royal Hospital. I might say: "Go to the hospital, and then turn right". The landmark referred to is the same as before, although in this case the expression omits the specifier. However, in this case the hospital is needed for an orientation task in a much smaller spatial environment. Hence, the elements activated in the cognitive representation of the listener by this expression are more detailed and closer to the hospital itself than in the first instance. For example, my friend will have no difficulties in orienting herself in the direction of the hospital at her current location, something that is not as easy for the taxi driver at Tullamarine, and also not necessary in his context.

Note that I could also have said to my friend: "At the corner is a hospital; turn right there", referring to an unspecific instance. In this case the hospital would serve only as a local landmark, since it can be distinguished only in vista space. When coming even closer to the hospital, one might refer no longer to the hospital as such, but to parts of it, such as the hospital's main entrance, or its Block B.

Symbolic and metaphoric use of landmarks in human communication is beyond the context of route or place descriptions, and hence, not considered here.

3.2 A rank order of landmarks

A cognitive rank order of landmarks appears natural, but for different reasons. We identify here orders in *prominence*, *uniqueness*, or *salience*.

3.2.1 Prominence

People's knowledge of landmarks is different, that means some are more prominent than others. For example, the Royal Hospital is known by more people than the (specific) 7-Eleven on the same corner as the hospital. Prominence imposes a partial order.

3.2.2 Uniqueness

Some landmarks are unique in larger spatial contexts than others. For example, the Royal Hospital is unique in Melbourne, but the 7-Eleven at the corner is unique only in the vista space of the intersection (as far as one being at that intersection can be sure). That means the size of the region which is associated with the use of a landmark in communication is a distinctive characteristics imposing another partial order. We call this region a reference region. These regions have no sharply defined boundaries. Note that even the visibility area of landmarks can vary largely, which will become relevant later (Section 5).

The order defined by uniqueness is independent from the order by prominence. In our example the hospital refers to different spatial contexts. Although the reference regions are different, the hospital's prominence is the same. In another example, Melbourne has a building of a unique blue color. Although the color is unique, the building is by far not as prominent as the Royal Hospital, in the context of route directions from the airport.

3.2.3 Salience

When searching for ways to automatically identify and select landmarks for route directions, some authors tried to identify observable and measurable properties of features that determine their salience [31, 7]. These attempts go back to a classification of Sorrows and Hirtle [34], who considered visible, structural and semantic properties of features as relevant in this context. A feature is salient, i.e., stands out out in its visible, structural or semantic properties, relative to other features nearby. A salient feature is called a landmark. Salience by these properties imposes another partial order, saying, for example, that the 7-Eleven at the corner is visually, semantically or structurally more salient than the parking garage next to it. However, this measure is dependent on prominence and uniqueness. A correlation with prominence can be concluded from the visible, structural or semantic properties of a landmark: at least one of them has made it prominent. A correlation with uniqueness can be concluded from the way of computation, which measures the distinction.

3.3 Properties of the rank order

Prominence, uniqueness, and salience lead in principle to different partial orders. Since salience measures are dependent on prominence as well as on uniqueness, we argue that salience measures reflect to some extent all three reasons for ranked landmarks in cognitive representations. Furthermore, the three approaches are all quantitative, although we have not discussed ways of quantification. This means, although the approaches allow a ranking of landmarks, the order is not yet a hierarchy of discrete levels.

It becomes now self-evident that features salient in a local neighborhood are local landmarks, and features salient in a global neighborhood are global landmarks. For example, the Royal Hospital in Melbourne is unique in Melbourne in its reputation, and in some of its highly specialized functions. It belongs to the salient features in Melbourne. In contrast, the 7-Eleven nearby is only salient when compared with other buildings nearby. Other than this correlation with their reference region, we have seen that the more salient landmarks are used over a range of contexts, and hence, have varying, context-dependent reference regions.

4 Formal models for landmark hierarchies

In this section we generate a levelled hierarchy of landmarks. For this purpose we adapt a method of identifying landmarks. Since being a landmark is a property relative to other features in the neighborhood, identification methods are regionalized. We exploit just this property for a levelled hierarchical ordering. We describe an operational procedure to come up with such an order.

4.1 Identification of landmarks

Some methods of identifying landmarks come with a continuous measure of salience [31, 11]. Continuous measures of salience can be used for classification, e.g., for selection processes [31]. Other methods immediately classify a feature, e.g., by the length of the description of the feature's distinctiveness [7], and thus, make a Boolean decision on salience. Boolean classification methods do not lend themselves automatically to create an order; one needs to consider additional properties for a quantification. We will adapt one of these regionalized classification methods further.

Other classification methods exist, which are not regionalized or have no spatial grounding at all. Winter and Tomko [47] explore a search engine's page counts of street addresses to generate a hierarchy of prominence. Geographic co-referencing is used to reconstruct spatial relations of nearness [29, 1], and within nearness relations one can again rank according to the direction of the references [38]. For example, an expression "the 7-Eleven near the Royal Hospital", and many similar expressions on the Web referring to the Royal Hospital, make this a prominent node. Since we want to exploit the duality of landmark and its reference region, these non-regionalized approaches to identify landmarks are not considered further in this paper.

4.2 Partitions by landmarks

We start from the assumption that any location in space can be described by references to landmarks. This means that any point in the Euclidean map plane is in at least one landmark's reference region, or, reference regions are jointly exhaustive. We also postulate that the reference regions are pairwise disjoint at each level of a hierarchy. This postulation is supported by two arguments: a practical and a theoretical one. From a practical perspective, if they were not pairwise disjoint the choice of landmarks to describe a location would be ambiguous. Since landmarks are used so successfully, this is implausible. From a theoretic perspective, one can assume a landmark forming a field of dominance decreasing with increasing distance from the landmark below any threshold. A landmark's reference region is then defined as the area where the dominance of the landmark is stronger than the dominance of any other competing landmark. This idea is formalized for cellular networks [18]. A similar concept of influence regions, represented by qualitative degrees of proximity, has been used for the definition of neighborhood [24].

These two properties mean that the reference regions of landmarks of similar salience can be considered to form partitions. For generating hierarchies, this insight can be exploited in two ways: (i) one can start from hierarchic partitions of space and search for the most salient feature in each cell, or (ii) one can start from identified landmarks and search for partitions.

In the first case, partitions could be made, for example, by a region quadtree, by resorting to administrative hierarchies, or by hierarchies of urban structure. However, a partition by a region quadtree has no cognitive motivation and would be translation variant. A partition by administrative hierarchies is not much more cognitively motivated since administrative boundaries are not well experienced by moving through urban space. A partition by urban structure has an arguable cognitive motivation, but no clear method for its generation. The duality between cells and landmarks remains questionable for this approach.

The other approach is to start with identifying landmarks. The regional identification methods discussed above (Section 4.1) are typically applied for the features in vista space of each street intersection in urban space. The reason is that an association of landmarks and street intersections is desired for the construction of route directions [15]. In route directions, landmarks are used to anchor the decisions and re-orientations along the route. We can assume that the set of vista spaces of all street intersections in an urban environment covers exhaustively the street space. This means that the identification procedure considers all features in the urban space.

After the identification of all landmarks at this local level we compute a Voronoi partition [25], with the landmarks as seeds. Note that at this level of density we cannot resort to vista spaces for partitioning. Landmarks can be in sight of each other, but still their reference regions, which will be bound to different street intersections, are not overlapping. In summary, we find strong cognitive motivation for this way of partitioning.

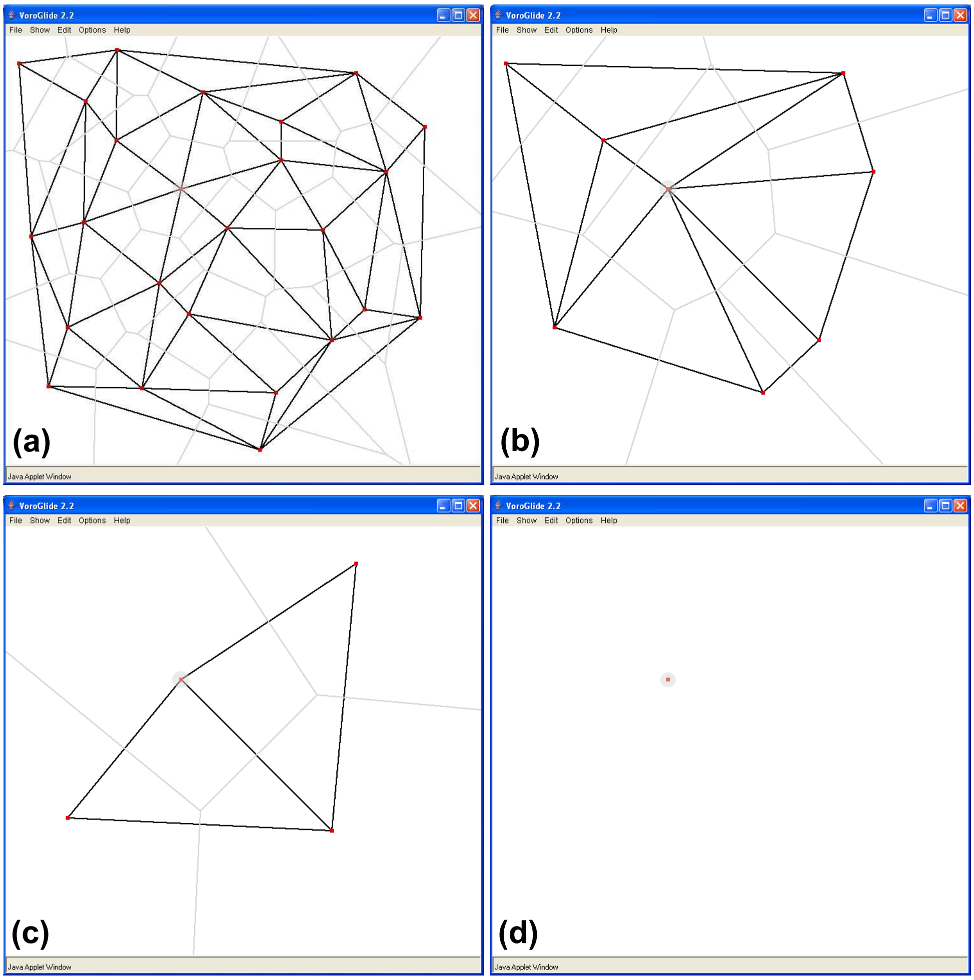

Figure 2: (a): Local landmarks with their reference regions (gray) and neighbored landmarks (black). (b)-(d): hierarchical levels of more and more salient landmarks. (d): the most salient landmark (marked down through all levels; note how its reference region changes).

The dual of the Voronoi partition, the Delaunay triangulation, links each landmark with all its neighbors. Figure 2a shows a set of local landmarks distributed over the map space, their Voronoi partition (gray), and the links to their immediate neighbors (black).

4.3 Generating a levelled hierarchy

A levelled hierarchy (Figure 2a-2d) can now be computed using a recursive procedure:

1. For each landmark l_k^i at level i, $l_k^i \in L^i$, compute the most salient landmark $l_{k.\max}$ in its immediate neighborhood $L_k^i = \{l_j \,|\, \text{dist}(l_j - l_k) \leq 1, l_j \in L^i\}$. The distance is defined as topological distance on the Delaunay triangulation.
2. Take the most salient landmarks $\{l_{k.\max} | l_k^i \in L^i\}$ as the set of landmarks $\{l_r^{i+1}\}$ at the next higher level of salience, L^{i+1}.
3. If there are more than one salient landmarks at level $i+1$, $|L^{i+1}| > 1$, compute the Voronoi partition and the Delaunay triangulation for all $l_k^{i+1} \in L^{i+1}$, and go back to step 1. Otherwise stop.

Note that the surviving landmarks at each new level existed already at the lower level, but their reference region changes. For example, study the reference region of the most salient landmark in Figure 2 through the levels a-d. Furthermore, in this levelled hierarchy, all landmarks already exist at the lowest level $i = 1$, and no new landmarks appear at higher levels. Landmarks can be characterized by the maximal level m that they appear in the hierarchy (l^m), and by the current level i in a discourse (l^i), with $i \leq m$. We call landmarks at level 1 *local* landmarks, and landmarks at levels $i > 1$ *global* landmarks. Further refinement of the terminology is possible.

A hierarchy generated in this way respects not only the partial order of the salience measures, it also forms a lattice. Its union is the most salient global landmark, and its intersection is the complete set of all point-like features of the space (with empty reference regions). A join and a meet exists for all pairs $\{l_i^m, l_j^n\}$ in L, defined by the links from neighborhoods to next-higher levels.

5 Realizing landmark hierarchies in an example

All known regionalized methods for landmark identification are based on an (arbitrary) choice of some neighborhood. But they are all developed in the realm of wayfinding and navigation, which means implicitly that they are interested in identifying local landmarks, or landmarks in local neighborhoods. A local neighborhood in the context of wayfinding and navigation can be defined by a potential decision point of a wayfinder in the urban environment (i.e., junctions), and all objects that intersect with its vista space. A local neighborhood has to contain more than one object. Identified objects are then objects that stand out among all objects in this neighborhood by visual, structural or semantic salience.

In this section we take one of the methods and apply it to different sizes of neighborhoods, starting with local neighborhoods. We expect that local landmarks stand out only in local neighborhoods, and global landmarks stand out in other sizes of neighborhoods as well. We will introduce the chosen method, present a test data set, and discuss the results of an hierarchical analysis according to Section 4.

5.1 Identifying landmarks by classification

The landmark identification procedure of Elias [8] automatically identifies building landmarks by classification. The approach consists of two different steps. In a first step, each potential decision point is investigated detecting the salient buildings in its vista space. This leads to a set of potential landmarks for each potential decision point. All objects in a set of potential landmarks are considered as equally salient for the potential decision point they are assigned to. In a second step, the number of potential landmarks is narrowed by introducing context specific characteristics of a chosen route, such as advance visibility while approaching the decision point, or the position of the landmark at the decision point. In our context, the second step is irrelevant, and hence, in the following experiment only the set of potential landmarks is generated.

The identification procedure investigates building data according to their geometric, semantic, and topological properties. These properties are modelled using different attributes assigned to the building objects, among them, for example, the function of the building, its distance to the road, and its orientation toward North. The objects are grouped into analysis units representing the vista space of each potential decision point using a 3D visibility analysis. For the landmark selection process, a measure from information theory is used, which is implemented in the machine learning algorithm *ID3* [30]. In order to use ID3 in an unsupervised way, an extension of the algorithm was developed. This modified ID3 algorithm is applied to the attribute database to identify automatically the most singular building objects in terms of their attribute values. Objects that differ considerably, in our case in one attribute value from all the others, are identified as a potential landmark, because they can be clearly referenced in the local neighborhood by naming this single characteristic. Examples would be *the red house* (color), *the large building* (size), or *the church* (function). The characteristics are not weighted, i.e., these examples are considered as equally salient—each of them standing out in one characteristic. As each building is evaluated, there can be several potential landmarks at one potential decision point.

5.2 Landmark test data area

The method was applied to data of a portion of the city of Hannover, Germany. For an area of about 5 km^2, the cadastral map, the topographic data set ATKIS, and a high-resolution 3D surface model generated from airborne laser scanning were available. The extent of the area is shown in Figure 3.

Figure 3: The test area, a part of the city of Hannover, Germany (approximately 4×4 km).

All junctions of the ATKIS road database were considered as potential decision points and used in the analysis process. Thereby 283 junctions are available as test data. In this area, there are 2200 buildings on the cadastral map given in form of polygons with additional thematic information. For each junction a 3D visibility analysis was conducted, using building polygons and the 3D surface model. For each junction, the building objects in its vista space establish its local neighborhood.

For each local neighborhood the data mining analysis process of the modified ID3 algorithm is applied. This way, for each junction a set of potential landmarks is identified. Altogether 868 building objects are chosen as potential landmarks, on average three for each junction. Because of the short distance between intersections in urban road networks, many buildings are visible from several junctions. Therefore, the analysis procedure can select individual buildings for more than one junction as a potential landmark. In the test area we find buildings that are selected only once, but also a building that is selected at 48 junctions as a potential landmark. Removing the redundant counts, the list of all potential landmarks in the test area consists only of 295 different building objects.

All potential landmarks are of equal salience. To determine a graduation of landmark salience, it is either necessary to take further route-specific aspects into account, such as advance visibility and quality, or the frequency of selection of a specific building is assessed as a measure for the global salience of the object. We choose the second approach for its independence from individual routes and better relation to the general experience of a city. In this way, visibility is ranked higher among all characteristics.

Table 1: Number of landmark objects in the seven hierarchy levels.

Level	0	1	2	3	4	5	6	7
Landmarks	295	95	32	11	4	3	2	1

5.3 Applying a Voronoi hierarchy onto the potential landmark selection

The introduced 295 different potential landmarks and their frequency of selection are used to establish a landmark Voronoi hierarchy. The potential landmark buildings are considered as Voronoi seeds in their own local neighborhood. In Figure 4 they are represented by a dot marking their centroid coordinates, and the number indicates their frequency counter. The Voronoi cells for these points are calculated as base level (level 0 in the hierarchy; Fig. 4 left).

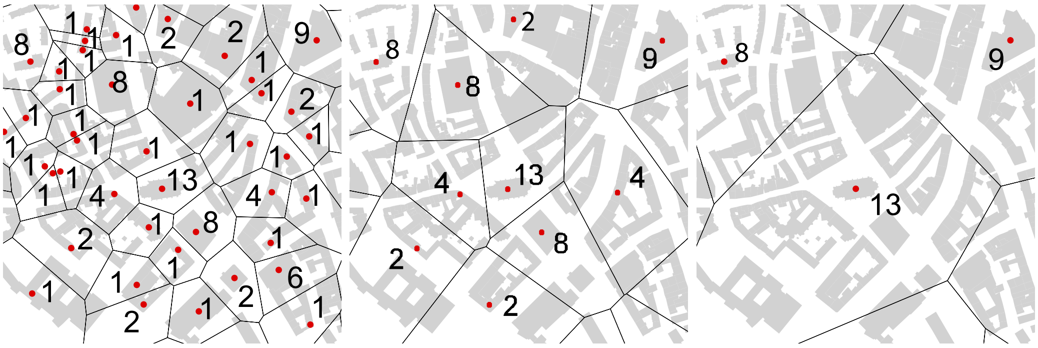

Figure 4: Voronoi cells with frequency counter. Left: base level, center: level 1, right: level 2.

So far, each potential landmark object is related to one Voronoi cell that represents the reference region of the landmark. For the next level in the hierarchy, the landmark selection is narrowed using the frequency of selection as a measure of salience, following the procedure introduced in Section 4.3. Step by step, each cell and its immediate neighbors are compared with each other, the cell object with the highest frequency counter is chosen as the most salient one and stored as an object of the next higher level (see Fig. 4, center and right). The procedure builds a new subset of landmarks with a Voronoi cell representing the reference region of the landmark at that level in the hierarchy. In this way, the number of objects decreases in each level. The process terminates when one single landmark—the most global landmark for the investigated area—is left.

In the test area seven hierarchy levels starting from the base level have to be processed (see Fig. 5 and Tab. 1). The single dot remained in level 7 represents the new town hall of Hannover, which is undoubtedly a prominent building in Hannover and one of its outstanding landmarks.

6 Evaluation of landmark hierarchies in context

Having derived the hierarchy of landmarks and their reference regions, the next question is how to make use of it. In this section we will investigate their use in *granular route directions*, which are route descriptions designed for locals. Granular route directions are route descriptions, or rather descriptions of route destinations, that make deliberate use of spatial hierarchies. In particular we will show how the duality between the landmarks and their reference regions allows for the evaluation of the spatial context of a route request. We will demonstrate how this duality is exploited in the selection of references to global or local landmarks.

6.1 Granular route directions with hierarchical spatial partitions

The construction of granular route directions based on hierarchical spatial partitions was previously explored in Tomko and Winter [44, 45]. In short, filter rules select only referents that uniquely and unambigu-

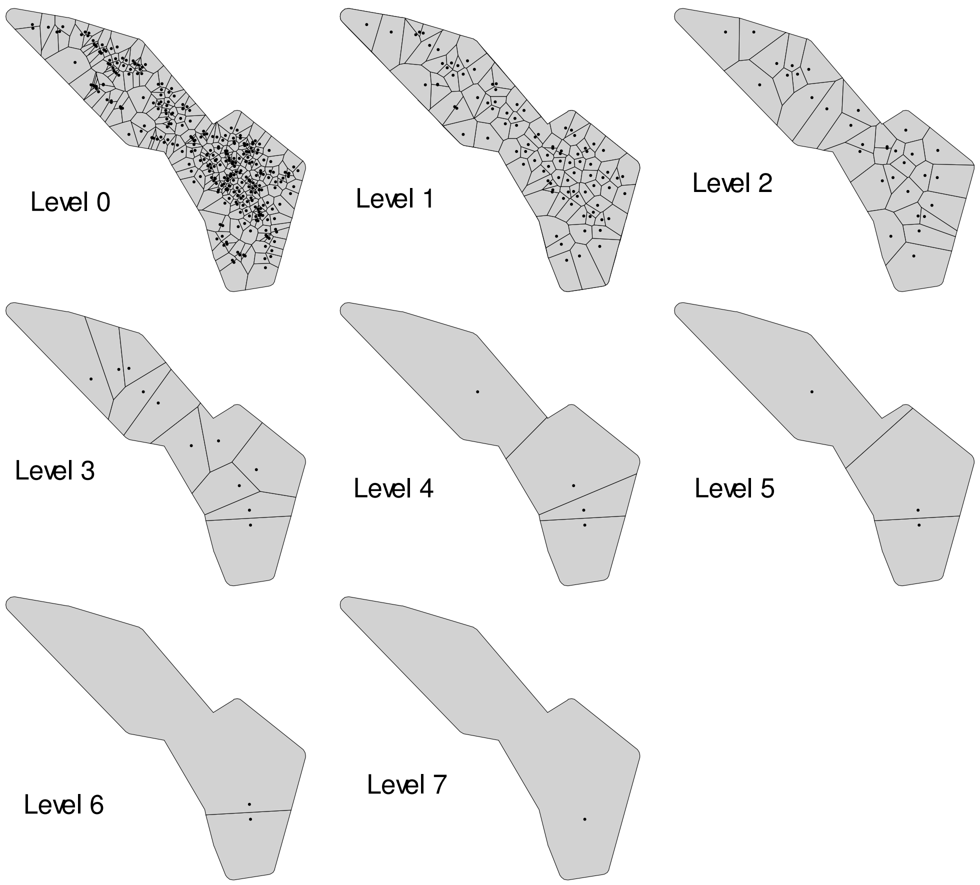

Figure 5: Base and higher hierarchy levels for the test area.

ously identify the destination of the route in a given context, i.e., for a given start location. In granular route directions, the localization of the destination is the goal of the directions, with the expectation that locals can find their way as soon as they have localized the destination.

This means the selection of the route as such is left to the wayfinder, and their tacit knowledge of the environment. This is quite different to other work on hierarchic or otherwise efficiently structured route directions [e.g., 41, 48, 12, 9]. Furthermore, every reference in granular route directions is evaluated against its information content, stating that it has to be relevant in the context of the route [35]. Altogether, granular route directions fulfill the requirement of being *referring expressions* [4], or minimal unambiguous descriptors of the referent, the route destination.

This requirement was accomplished by a recursive application of a few topological rules on hierarchical spatial partitions [44, 45]. The initial step is searching for the highest hierarchy level in which the current location and destination are still in a topological distance larger than 1. As the wayfinder progresses towards the destination, the location changes and consecutive references are generated with consideration of this changed context. The rules are:

1. If current location or destination are not members of the hierarchy, throw exception.
2. If current location and destination are identical, stop.
3. If current location and destination are neighbors, stop and switch to turn-based directions.
4. If current location and destination have identical or neighboring direct superordinate elements, return a reference to the destination and stop.
5. If an element is a common ancestor of current location and destination, move down a level.
6. If an element is neighbor with an ancestor element of the current location, move down a level.
7. Otherwise: return a reference and continue.

Next we will examine these rules on the landmark hierarchy, with the working hypothesis that the same rules applied to the hierarchy of reference regions lead to referring expressions based on landmarks.

6.2 Granular route directions with landmark hierarchies

Neighborhood is the fundamental topological relation driving the selection of references in granular route directions. Two areas are neighbors if they share a common boundary. Translated into our context, we say that two landmarks are neighbors if their respective reference regions within the same hierarchical level are adjacent. Since the reference regions are modelled as Voronoi polygons, neighborhood is represented by their dual, the Delaunay triangulation. Relations between the levels of the hierarchy are retained in the list of associations from the ordering process of landmarks.

So, can we apply the above rules for the construction of granular route directions to landmarks? The rules were originally developed for tree-like hierarchical partitions. As we have seen, the landmark hierarchy does not form a tree. Furthermore, some landmarks are their own superordinate elements, due to the peculiarities of the generation of the hierarchy. Hence, the rules need an adaptation to be able to navigate in a hierarchy with multiple inheritances and multiple appearances of the same landmark.

Rules 1–4 apply to both hierarchies equally. The application of the Rules 5 and 6 to the hierarchy of landmarks has, however, several shortcomings:

- The reference is always made to the landmark, and not to its reference area. The identity of the landmark is unique even if its reference area changes across levels. Thus, it is possible that a reference is made to a landmark that at a higher hierarchical level is an ancestor of both the current location and the destination. Rule 5 is therefore irrelevant.

- It is impossible to test the neighborhood relation between the reference regions of landmarks with the distance of hierarchical levels greater than 1. If the difference is one hierarchical level, this test can be done partially through the association information. Beyond this, there is no conceptual grounding of the neighborhood relation between point-like landmarks, even if spatial analysis can be done between the respective reference regions.

A restricted set of rules, dealing with these shortcomings, produces granular route directions with excessive information, mostly in the form of redundant references to the same landmark. Furthermore, the algorithm is still not ready to deal with multiple ancestors of an element. This means that, at any level of the hierarchy, the set of possible referents is not restricted to one landmark.

Hence, we propose a further refined set of rules, enriched by mechanisms to deal with redundant referents in the resulting route directions, as well as with the selection among multiple possible referents. We propose two new rules for the construction of landmark-based granular route directions:

1. Among possible referents, priority is given to the referents along the route. If multiple landmarks satisfy this condition, the landmark closest to the destination is selected.
2. If a landmark is referred to multiple times, remove all but one reference.

This means basically that in the context of landmark-based granular route directions topological distance is not the only rule that may apply. The consideration of the route context is a plausible mechanism to select among multiple referents, especially at finer levels of granularity.

In our implementation, a route is specified as a set of local reference regions. This route is never communicated to the wayfinder. It provides, however, the spatial context for the identification of consecutive references.

6.3 Test of landmark-based granular route directions

The modified rules were implemented in the functional programming language Haskell, and applied to the example of the City of Hannover (Section 5). We will describe the process of identification of the referents for a route from the University of Hannover to the State Opera (Fig. 6). This route was chosen from a map and intersected with the Voronoi partition. For clarity and brevity, we restrict our example to hierarchy levels 2 and above, abandoning the large set of local landmarks.

In the figure, the university at level 2 is labelled h2_H097TLK, and the destination, the opera, is labelled h2_H01FM8E. The complete route at this level 2 is defined by a sequence of reference regions:

```
route = [h2_H097TLK, h2_H063YJC, h2_H074YH2, h2_H03PO3Z,
    h2_H04SBR1, h2_H01BHXG, h2_H01P2HN, h2_H04PTS0, h2_H01FM8E]
```

This route cuts across the city and crosses the reference regions of several global landmarks. It does not mean that the wayfinders reach those landmarks, but merely that while proceeding towards the destination, they enter their reference regions. Applying the rules for granular route directions generate the following description:

```
directions = [H06Y0NB, H03NG5F, H04PTS0, H01FM8E]
```

Note the missing prefixes for hierarchical levels. The reference is made to the landmark as such, not to its specific reference regions. In this case, the destination is found in the proximity of the first landmark specified, the global landmark H06Y0NB (Hannover town hall). The context of the route is restricted to this area, and a consecutive reference of finer granularity is provided (H03NG5F). In this manner, we proceed from general references (landmarks with reference regions covering major parts of the city) to more and more local landmarks. References provided may have no spatial overlap with the route as such, but they provide a description of the destination from the approaching direction of the wayfinder. The two consecutive references H03NG5F and H04PTS0 are filled in at consecutive steps of the algorithm through

Figure 6: A route (highlighted in gray) between the University of Hannover (h2_H097TLK) and the State Opera (h2_H01FM8E), composed of level 2 reference regions. Referents identified for granular route directions are labelled.

the same process (Fig. 7). The route description is completed with the reference to the destination itself (H01FM8E).

We can interpret these results as the following route description generated by a human:

A:"Where is the show this evening?"
B:"It is in the State Opera, the building next to the Kröpcke, not far from the town hall and the Maritim-Hotel building."

In this example we have made additional use of spatial context. Every point within a reference region is considered to be within the spatial context of the respective landmark, and is thus *next to* it. This is close to the tectonic plates hypothesis [3]. The association of arbitrary points in space with their reference landmarks can change between consecutive granularity levels.

The result is a plausible granular route direction from the university to the opera. Consultation with a local

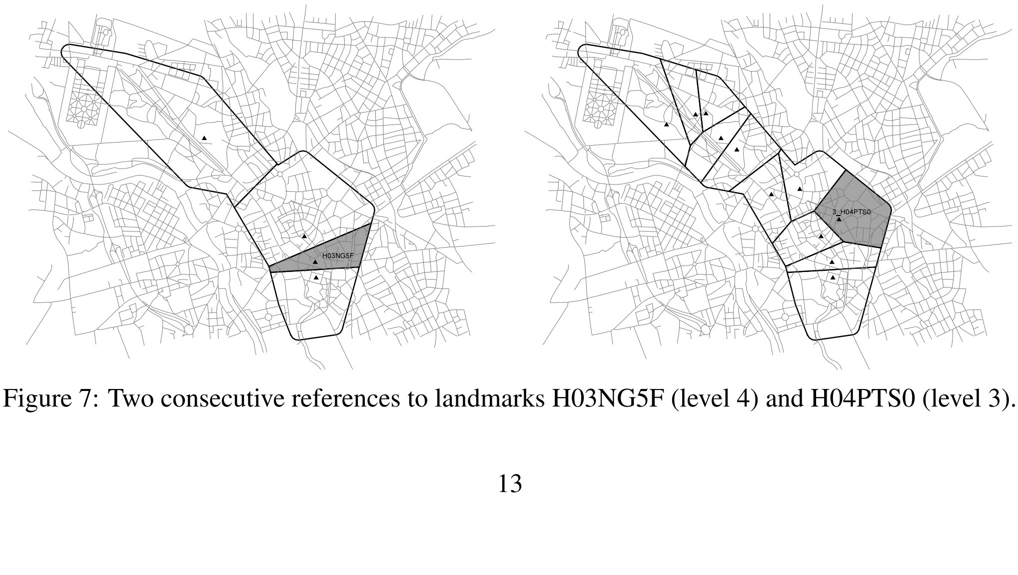

Figure 7: Two consecutive references to landmarks H03NG5F (level 4) and H04PTS0 (level 3).

expert shows that, depending on the spatial context of the chosen route, our destination description presents referents likely to be chosen by a local. Only the reference to the Maritim-Hotel was felt to be excessive. The reference was made due to the high visual salience of the hotel, which in this case seems not to match with prominence. This observation indicates also that the distinctions of salience in the top levels of the hierarchy might be not pragmatic.

Contexts other than spatial context, such as mode of transport, make a difference, but are not considered in this paper. In the case of a pedestrian, for example, a local would very likely refer directly to the opera, as the building as such is prominent for every local. In the case of a car driver, though, the route from the university to the opera is lengthy and complex (the opera is in a pedestrian zone). Hence, the references provided by our model show a granularity likely to be used by locals for car drivers.

Let us consider also the reverse route from the opera to the Institute for Chemistry on the university campus. It is one of the buildings adjacent to the very prominent main building of the university. The resulting directions consist of only two referents: the main building and the building of the Institute for Chemistry. This result again was assessed by a local expert as a plausible result.

The simple combination of the topological distance and the context of the route provided a set of rules enabling the construction of granular route directions. Based on a hierarchical ranking of landmarks, the results provide plausible directions for locals. Topology and the route context can be further enhanced by other considerations, to provide a better match with those provided by humans in specific situations.

7 Conclusions

The ultimate goal of our research is to design spatial information services for everyday decisions that better reflect the way people communicate about spatial information. We believe that services reflecting human spatial communication are more easily accepted, because they require less learning, lower cognitive workload, and lower memory load (Tomko and Winter [43], for example, argue that human granular route directions are designed to be kept in short term memory). In this way service users should find spatial descriptions more intuitive, their comprehension easier, and hence, feel more satisfied or more comfortable in using these descriptions as oppose to descriptions from actual services.

In this paper we have taken an engineering approach to reflect in a formal model a process known from human spatial cognition: the hierarchical representation, reasoning and communication of spatial information. In particular we focus on formal models for landmarks and their use in route or place descriptions. These models enable us to use geographical databases to derive route descriptions that are close to human ones, introducing from context concepts such as *near to*. In these models spatial context is the key for the selection of references.

We have presented a model to build hierarchies of landmarks from salience, extending one of the models to automatically identify landmarks. With a ranking based on visibility, we even found a cognitive relevant ranking criterion. This model was tested for a part of the city of Hannover, Germany, and showed reasonable behavior. We have then taken a model to generate granular route directions from hierarchical partitions, and modified it for the particularities of the landmark hierarchies. This model was tested on the same dataset and also showed reasonable behavior. Hence, we could prove our hypothesis.

Refinements and more extensive tests of these models are possible. For example, ranking based on visibility did not take into account the distance, or the proportions in the field of view, and ranking factors other than visibility could also be investigated. Any change of the ranking method influences the resulting hierarchy. Another example is the method to identify the basic set of landmarks. It currently applies equal weights to the different salience factors; a different weighting would lead to different choices of landmarks, and hence, different hierarchies. Equally, the rules on landmark hierarchies for granular route directions can be refined and extended. The current set of rules was tested on the dataset of Hannover, but with more datasets one could find special cases not yet considered.

Another refinement concerns the partition of reference regions. Especially in urban space, taking barriers into account should constrain the partitioning process. This includes the preservation of boundaries across hierarchical levels.

Having solved the problem of hierarchically ordering landmarks is an essential condition for hierarchic route or place descriptions. However, landmarks are only one type of element relevant for these descriptions. Other elements of the city, such as places, paths, barriers, and districts [21], are also experienced by people, form parts of their cognitive spatial representations, and hence, have hierarchical cognitive order as well. While some work on the other elements is already done [45] or underway, this article contributes the hierarchies of landmarks. An open question for future work is how to link these hierarchies in the generation of route or place descriptions.

We also expect more subtle links between different hierarchies. For example, in the current approach all landmarks are considered point-like and of the same type (building), and accordingly, all appear at the same local level. If we allow for other types of landmarks, then landmarks can appear at higher levels for the first time, a property of granularity. For example, the site of the world exhibition in Hannover (EXPO2000) might be considered as a landmark, but consists of multiple buildings. Hence, a landmark *EXPO2000* does not exist at the local level of buildings. Therefore, the integration of landmarks with other types of elements of the city is an interesting remaining question. Relationships to scale space [16, 19] and to map scale hierarchies [39] are expected.

Acknowledgements

Geodata by courtesy of National Mapping Agency of Lower Saxony, Germany (LGN).

References

[1] Arampatzis, A., van Kreveld, M., Reinbacher, I., Jones, C. B., Vaid, S., Clough, P., Joho, H., Sanderson, M., 2006. Web-Based Delineation of Imprecise Regions. Computers, Environment and Urban Systems in press.

[2] Bateman, J., 1999. Using Aggregation for Selecting Content When Generating Referring Expressions. In: Proceedings of the 37th Annual Meeting of the Association for Computational Linguistics on Computational Linguistics. ACL, College Park, Maryland, pp. 127–134.

[3] Couclelis, H., Golledge, R. G., Tobler, W., 1987. Exploring the Anchorpoint Hypothesis of Spatial Cognition. Journal of Environmental Psychology 7, 99–122.

[4] Dale, R., 1992. Generating Referring Expressions: Constructing Descriptions in a Domain of Objects and Processes. MIT Press, Cambridge, Mass.

[5] Dale, R., Reiter, E., 1995. Computational Interpretations of the Gricean Maxims in the Generation of Referring Expressions. Cognitive Science 19 (2), 233–263.

[6] Denis, M., 1997. The description of routes: A cognitive approach to the production of spatial discourse. Current Psychology of Cognition 16 (4), 409–458.

[7] Elias, B., 2003. Extracting Landmarks with Data Mining Methods. In: Kuhn, W., Worboys, M. F., Timpf, S. (Eds.), Spatial Information Theory. Vol. 2825 of Lecture Notes in Computer Science. Springer, Berlin, pp. 398–412.

[8] Elias, B., 2006. Extraktion von Landmarken für die Navigation. Dissertation, University of Hannover, wissenschaftliche Arbeiten der Fachrichtung Geodäsie und Geoinformatik der Universität Hannover, Nr. 260.

[9] Elias, B., Sester, M., 2006. Incorporating landmarks with quality measures in routing procedures. In: Raubal, M., Miller, H. J., Frank, A. U., Goodchild, M. F. (Eds.), Geographic Information Science. Vol. 4197 of Lecture Notes in Computer Science. Springer, Berlin, pp. 65–80.

[10] Frank, A. U., Timpf, S., 1994. Multiple Representations for Cartographic Objects in a Multi-Scale Tree - An Intelligent Graphical Zoom. Computers and Graphics 18 (6), 823–829.

[11] Galler, I., 2002. Identifikation von Landmarks in 3D-Stadtmodellen. Diploma thesis, Rheinische Friedrich-Wilhelms-Universitt Bonn.

[12] Hansen, S., Richter, K.-F., Klippel, A., 2006. Landmarks in OpenLS: A Data Structure for Cognitive Ergonomic Route Directions. In: Raubal, M., Miller, H. J., Frank, A. U., Goodchild, M. F. (Eds.), Geographic Information Science. Vol. 4197 of Lecture Notes in Computer Science. Springer, Berlin, pp. 128–144.

[13] Hirtle, S. C., Jonides, J., 1985. Evidence of Hierarchies in Cognitive Maps. Memory and Cognition 13 (3), 208–217.

[14] Jarvella, R. J., Klein, W. (Eds.), 1982. Speech, Place, and Action. John Wiley & Sons, Chichester, NY.

[15] Klippel, A., Winter, S., 2005. Structural Salience of Landmarks for Route Directions. In: Cohn, A. G., Mark, D. M. (Eds.), Spatial Information Theory. Vol. 3693 of Lecture Notes in Computer Science. Springer, Berlin, pp. 347–362.

[16] Koenderink, J. J., 1984. The Structure of Images. Biological Cybernetics 50 (5), 363–370.

[17] Kuipers, B., 1998. A Hierarchy of Qualitative Representations for Space. In: Freksa, C., Habel, C., Wender, K. F. (Eds.), Spatial Cognition. Vol. 1404 of Lecture Notes in Artificial Intelligence. Springer, Berlin, pp. 337–350.

[18] Lang, D., Winter, S., Frank, A. U., 2001. Neighborhood Relations between Fields with Applications to Cellular Networks. GeoInformatica 5 (2), 127–144.

[19] Lindeberg, T., 1994. Scale-Space Theory in Computer Vision. Kluwer Academic Publishers, Dordrecht.

[20] Lovelace, K. L., Hegarty, M., Montello, D. R., 1999. Elements of Good Route Directions in Familiar and Unfamiliar Environments. In: Freksa, C., Mark, D. M. (Eds.), Spatial Information Theory. Vol. 1661 of Lecture Notes in Computer Science. Springer, Berlin, pp. 65–82.

[21] Lynch, K., 1960. The Image of the City. MIT Press, Cambridge.

[22] Montello, D. R., 1993. Scale and Multiple Psychologies of Space. In: Frank, A. U., Campari, I. (Eds.), Spatial Information Theory. Vol. 716 of Lecture Notes in Computer Science. Springer, Berlin, pp. 312–321.

[23] Montello, D. R., Goodchild, M. F., Gottsegen, J., Fohl, P., 2003. Where's Downtown? Behavioral Methods for Determining Referents of Vague Spatial Queries. Spatial Cognition and Computation 3 (2&3), 185–204.

[24] Moulin, B., Kettani, D., 1999. Route generation and description using the notions of object's influence area and spatial conceptual map. Spatial Cognition and Computation 1 (3), 227–259.

[25] Okabe, A., Boots, B., Sugihara, K., Chiu, S. N., 2000. Spatial Tesselations: Concepts and Applications of Voronoi Diagrams, second edition Edition. John Wiley & Sons, Chichester, United Kingdom.

[26] Paraboni, I., Deemter, K. v., 2002. Generating Easy References: The Case of Document Deixis. In: Second International Conference on Natural Language Generation (INLG 2002). New York, USA, pp. 113–119.

[27] Plumert, J. M., Carswell, C., DeVet, K., Ihrig, D., 1995. The Content and Organization of Communication about Object Locations. Journal of Memory and Language 34, 477–498.

[28] Presson, C. C., Montello, D. R., 1988. Points of Reference in Spatial Cognition: Stalking the Elusive Landmark. British Journal of Developmental Psychology 6, 378–381.

[29] Purves, R. S., Clough, P., Joho, H., 2005. Identifying Imprecise Regions for Geographic Information Retrieval Using the Web. In: Billen, R., Drummond, J. E., Forrest, D., Joo, E. (Eds.), GIS Research UK. Glasgow, UK, pp. 313–318.

[30] Quinlan, J. R., 1986. Induction of Decision Trees. Machine Learning 1 (1), 81–106.

[31] Raubal, M., Winter, S., 2002. Enriching Wayfinding Instructions with Local Landmarks. In: Egenhofer, M. J., Mark, D. M. (Eds.), Geographic Information Science. Vol. 2478 of Lecture Notes in Computer Science. Springer, Berlin, pp. 243–259.

[32] Schelling, T. C., 1960. The Strategy of Conflict. Harvard University Press, Cambridge, Mass.

[33] Shanon, B., 1983. Answers to Where-Questions. Discourse Processes 6, 319–352.

[34] Sorrows, M. E., Hirtle, S. C., 1999. The Nature of Landmarks for Real and Electronic Spaces. In: Freksa, C., Mark, D. M. (Eds.), Spatial Information Theory. Vol. 1661 of Lecture Notes in Computer Science. Springer, Berlin, pp. 37–50.

[35] Sperber, D., Wilson, D., 1986. Relevance – Communication and Cognition. Basil Blackwell, Oxford.

[36] Stevens, A., Coupe, P., 1978. Distortions in Judged Spatial Relations. Cognitive Psychology 10 (4), 422–437.

[37] Strohecker, C., 2000. Cognitive Zoom: From Object to Path and Back Again. In: Freksa, C., Brauer, W., Habel, C., Wender, K. F. (Eds.), Spatial Cognition II. Vol. 1849 of Lecture Notes in Artificial Intelligence. Springer, Berlin, pp. 1–15.

[38] Tezuka, T., Yokota, Y., Iwaihara, M., Tanaka, K., 2004. Extraction of Cognitively-Significant Place Names and Regions from Web-Based Physical Proximity Co-occurrences. In: Zhou, X., Su, S., Papazoglou, M. P., Orlowska, M. E., Jeffrey, K. G. (Eds.), Web Information Systems (WISE 2004). Vol. 3306 of Lecture Notes in Computer Science. Springer, Berlin, pp. 113–124.

[39] Timpf, S., 1998. Hierarchical Structures in Map Series. Ph.d. thesis, Technical University Vienna.

[40] Timpf, S., Frank, A. U., 1997. Using Hierarchical Spatial Data Structures for Hierarchical Spatial Reasoning. In: Hirtle, S. C., Frank, A. U. (Eds.), Spatial Information Theory (COSIT ’97). Vol. 1329 of Lecture Notes in Computer Science. Springer, Berlin, pp. 69–83.

[41] Timpf, S., Kuhn, W., 2003. Granularity Transformations in Wayfinding. In: Freksa, C., Brauer, W., Habel, C., Wender, K. F. (Eds.), Spatial Cognition III. Vol. 2685 of Lecture Notes in Artificial Intelligence. Springer, Heidelberg, pp. 77 – 88.

[42] Timpf, S., Volta, G. S., Pollock, D. W., Frank, A. U., Egenhofer, M. J., 1992. A Conceptual Model of Wayfinding Using Multiple Levels of Abstraction. In: Frank, A. U., Campari, I., Formentini, U. (Eds.), Theories and Methods of Spatio-Temporal Reasoning in Geographic Space. Vol. 639 of Lecture Notes in Computer Science. Springer, Berlin, pp. 348–367.

[43] Tomko, M., Winter, S., 2006. Considerations for Efficient Communication of Route Directions. In: Cartwright, W., Yoshida, H., Andrienko, G. (Eds.), ISPRS Workshop on Spatial Data Communication and Visualization. ISPRS, Vienna, Austria.

[44] Tomko, M., Winter, S., 2006. Initial Entity Identification for Granular Route Directions. In: Kainz, W., Riedl, A., Elmes, G. (Eds.), Progress in Spatial Data Handling. Springer, Vienna, Austria, pp. 43–60.

[45] Tomko, M., Winter, S., 2006. Recursive Construction of Granular Route Directions. Journal of Spatial Science 51 (1), 101–115.

[46] Tsuchiya, P. F., 1988. The Landmark Hierarchy: A New Hierarchy for Routing in Very Large Networks. ACM SIGCOMM Computer Communication Review 18 (4), 35–42.

[47] Winter, S., Tomko, M., 2006. Translating the Web Semantics of Georeferences. In: Taniar, D., Rahayu, J. W. (Eds.), Web Semantics & Ontology. Idea Publishing, Hershey, PA, pp. 297–333.

[48] Wuersch, M., Caduff, D., 2005. Refined Route Instructions Using Topological Stages of Closeness. In: Li, K.-J., Vangenot, C. (Eds.), Web and Wireless Geographical Information Systems. Vol. 3833 of Lecture Notes in Computer Science. Springer, Berlin, pp. 31–41.