Prospects and Challenges of Landmarks in Navigation Services

Kai-Florian Richter

Abstract In the past decades, empirical research has established the importance of landmarks in our understanding of and communication about space. These findings have led to the development of several computational approaches for the automatic identification and integration of landmarks in navigation instructions. However, so far this research has failed to make any impact on commercial services. This chapter will discuss reasons for this failure. It will develop a categorization of existing approaches and highlight their shortcomings. Finally, principles and methods of user-generated content will be identified as a promising, feasible way forward to future landmark-based navigation services.

Keywords Landmarks · User-generated content · Route directions · Location-based services

1 Introduction

In research on people's understanding of space, landmarks have been consistently shown to be of great importance, going back at least to Lynch's seminal work on 'the image of the city' (1960) which looked at how long-term residents conceptualize their cities' layout and social structure. Landmarks are important in learning environments (Siegel and White 1975) and in forming mental representations of environments (Couclelis et al. 1987; Hirtle and Jonides 1985). When communicating about an environment, for example, as when giving route

K.-F. Richter (✉)
Department of Infrastructure Engineering, The University of Melbourne, Melbourne, Australia
e-mail: krichter@unimelb.edu.au

M. Raubal et al. (eds.), *Cognitive and Linguistic Aspects of Geographic Space*, Lecture Notes in Geoinformation and Cartography, DOI: 10.1007/978-3-642-34359-9_5,

directions, people use landmarks to anchor actions in space or to provide confirmation that the right track is still being followed (Denis 1997; Lovelace et al. 1999; Michon and Denis 2001).

Not surprisingly, landmarks are highly desired additions to automatic navigation services, such as car navigation systems. They are a top feature request of users (May et al. 2003). Using prototypical research systems, landmarks were found to improve users' performance and satisfaction with such systems in both car navigation (Burnett et al. 2001; Ross and Burnett 2001) and pedestrian wayfinding (May et al. 2003; Ross et al. 2004). However, they have hardly found their way into commercial systems and services—with a notable exception of the Australian routing service *WhereiS*.[1]

This chapter will explore some of the reasons why landmarks fail in end-user products. It will do so by analyzing and categorizing existing approaches for the identification and integration of landmarks in wayfinding instructions, and then pointing out shortcomings and challenges of these approaches. While there have been such analyses before (Sadeghian and Kantardzic 2008; Tezuka and Tanaka 2005), they restricted themselves to the extraction (identification) of landmarks; also, they miss some important recent developments in the field. The chapter will also propose novel ways of including landmarks that employ mechanisms of user-generated content and Web 2.0 technology. But first of all it will explain what is meant by the term 'landmark.'

1.1 What is a Landmark?

Lynch (1960) defined a landmark to be a readily identifiable object which serves as external reference point. This definition is frequently picked up in the literature, often resulting in landmarks being conceived as point-like features along a route. However, anything that sticks out from the background may serve as a landmark (Presson and Montello 1988). In light of this broad definition, the Urban Knowledge Data Structure (Hansen et al. 2006; Klippel et al. 2009) provides an elaborate formal specification of which types of geographic features may serve as landmarks in automatically generated route directions, from signage found along a street and individual buildings, such as churches, to linear features, such as rivers (Richter 2007), to salient street intersections, such as roundabouts (Klippel et al. 2005).

The "sticking out from the background" of a feature is often defined through its salience (Elias 2003; Raubal and Winter 2002). Sorrows and Hirtle (1999) identified three key characteristics of landmarks that influence this salience: (1) singularity, i.e., contrast with surroundings; (2) prominence of spatial location; (3) content, i.e., meaning or cultural significance. Several approaches aim at covering

[1] http://www.whereis.com.au

these characteristics in calculating salience values for landmark candidates. These will be discussed in the next section.

1.2 Landmark Identification and Landmark Integration

The inclusion of landmark references in automatically generated navigation instructions requires two steps: (1) the identification of features that may serve as landmarks in principle. In the following, these features will be referred to as *landmark candidates*; (2) the selection of some of these candidates to be included in the instructions. Often, these two steps are seen as independent (Elias et al. 2005). Consequently, most approaches that will be presented in the following either focus on the identification of landmark candidates or on the selection of features from a set of landmark candidates that are then integrated into the generated instructions.

1.2.1 Landmark Identification

In general, landmark identification is performed by specifying a region in which landmarks are to be sought (e.g., an area around an intersection), and then identifying outliers relative to other features in this region, i.e., finding salient features (Sadeghian and Kantardzic 2008).

The first approach to the automatic identification of landmarks was presented by Raubal and Winter (2002). It inspired several extensions and further approaches. Their approach reflects the three landmark characteristics of Sorrows and Hirtle (1999) by taking into account different attributes of building façades (e.g., area, color, signs, visibility) in a weighted sum for calculating individual buildings' salience. These façades serve as point-like landmarks along a route; the required data is supposed to be stored in a spatial database (GIS). Employing a user survey, weights for the individual parameters were set for specific situations, accommodating for differences between day and night (Winter et al. 2005). New situations require an adaptation of these weights, which likely requires new user studies.

In Winter (2003), calculation of salience accounts for advance visibility, i.e., how soon and for how long a façade is visible when considering the direction of travel. This is further refined in Klippel and Winter (2005), where locations of landmark candidates along a route are taken into account—termed structural salience by the authors. The location of a landmark influences the ease of conceptualizing turning actions and, thus, determines the ease of understanding instructions (Richter and Klippel 2007).

A similar approach to Raubal and Winter (2002) is taken by Elias (2003). She uses machine learning techniques to identify the most salient objects in a spatial data set. These objects are considered to be point-like entities. Winter et al. (2008) combine the two approaches of Raubal and Winter and of Elias to construct a hierarchy of landmarks based on each individual candidate's salience. This hierarchy is used in the generation of destination descriptions (Tomko and Winter 2009), which, accordingly, then is an example for integrating landmarks into wayfinding instructions.

Others have explored data mining approaches to identify landmark candidates for navigation services. Tomko (2004) used requests to Internet search engines looking for street names and subsequent filtering mechanisms to identify potential salient features (buildings) along a previously calculated route. Both search and filtering of results were done manually in this case, but may be automated with specifically tailored web services. In Tezuka and Tanaka (2005), text mining methods are used to mine WWW documents in order to identify prominent point-like features; prominence is based on how authors of these documents refer to the features. Mining of landmark information is further discussed below.

1.2.2 Landmark Integration

Caduff and Timpf (2005) presented an algorithm that calculates a route through a network based on the presence of point-like landmarks at decision points (nodes). It tries to navigate a wayfinder along a route that has a landmark at every decision point. They did not specify how these landmarks are identified, but rather assume their existence. The same holds for the approach by Richter (Richter 2007; Richter 2008; Richter and Klippel 2007), which integrates landmarks into an abstract specification of route directions that follow cognitive principles of direction giving. His approach selects those landmarks from a set of landmark candidates that are best suited to describe actions to be performed (cf. Klippel and Winter 2005). In a similar line, Elias and Sester (2006) used a modified Dijkstra shortest path algorithm to find a route through a network that integrates landmarks. Weights in the network are adapted according to the permanence, visibility, usefulness of location, uniqueness, and brevity of description of landmarks. Landmarks are assumed to be point-like (buildings) and are determined using the approach by Elias (2003).

The CORAL system by Dale et al. (2005) produces natural language instructions for route following, mimicking human principles of direction giving. Integration of landmarks is based on work by Williams (1998), which employs common-sense rules for selecting landmarks in indoor environments. The approach is not well documented, but seems to use location of a landmark and travel direction as parameters.

Recently, Duckham et al. (2010) explored using categories of features instead of their individual properties to determine suitability as a landmark. They combined a category's general suitability, its uniqueness in an area and a feature's location along a route to select those features best suited to describe how to follow the given route. This approach is implemented in the *WhereiS* route service using categories taken from the yellow pages. A similar approach was taken by Wagner (2009).

2 Landmarks in Navigation Services: A Categorization

This section develops a categorization of the approaches presented in the previous two sections. A first, broad categorization is already done there: the distinction between landmark identification and landmark integration. This distinction is the top level of the proposed categorization. Further, some approaches focus on properties of the features themselves, i.e., how they differ from other features in their surrounding. Other approaches account for the location of landmark candidates along a route to assess their suitability as references in instructions. The former is a static view on landmarks, the latter a dynamic view. This difference in views is similar to the distinction between *structure* and *function* in wayfinding as introduced by Klippel (Klippel 2003).

The (assumed) source of data that landmark identification is based on also differs between the approaches. Some use spatial databases of the kind attached to a typical geographic information system (GIS), some use data from the web (general websites; or web catalogs, such as yellow pages), and others do not specify their data source (marked as *abstract* in Table 1). Finally, approaches differ in the conceptual geometry of landmark candidates [points or more complex features, i.e., polygons; Hansen (2006)] and in whether they aim to identify individual features (instances) or categories of features (types).

Table 1 shows a matrix that categorizes the approaches presented in the previous two sections according to these criteria.

The first observation to make when looking at this matrix is that approaches to landmark identification predominantly work on a structural level, while landmark integration is on a functional level. This supports the statement made previously that identification of landmark candidates and selection of landmarks to be integrated in wayfinding instructions are considered to be independent steps. Landmark identification needs to find all features that may serve as a landmark in principle, i.e., are sufficiently salient. For a given data set, this may be done in a preprocessing step. The resulting set of landmark candidates then can be used as a pool of potential landmarks to select from when generating route directions for any route through the environment. Further, salience is a local feature (Elias 2003), in that a feature needs to stick out from its neighboring features. Thus, for landmark identification this static view on features' properties is useful as it allows determining a feature's general suitability. The two approaches that (also) are on a functional level already assume a specific given route for which landmarks are to be identified. They also partly cover the integration step, particularly Klippel and Winter (2005).

Landmark integration, on the other hand, needs to ensure that the referenced landmarks are actually useful in a navigation context. Landmarks need to be visible in the direction of travel, sensibly describe a (turning) action, and support conceptualization of the instructions. These characteristics are functional, as they depend on the specific route at hand. In the integration step, the landmarks chosen are not necessarily the most salient landmarks, but those that are most relevant for

Table 1 Matrix of approaches to landmark identification and integration

	Identification													
	Structure							Function						
	Data Source			Geometry				Data Source			Geometry			
	GIS	WWW	Abstract	Point	Polygon	Individual	Category	GIS	WWW	Abstract	Point	Polygon	Individual	Category
Elias (2003)	■			■		■								
Klippel and Winter (2005)										■	■		■	
Raubal and Winter (2002)	■			■		■								
Tezuka and Tanaka (2005)		■		■		■								
Tomko (2004)		■		■		■								
Winter (2003)	■			■		■		■			■		■	
Winter et al. (2008)	■			■		■								

	Integration													
	Structure							Function						
	Data Source			Geometry				Data Source			Geometry			
	GIS	WWW	Abstract	Point	Polygon	Individual	Category	GIS	WWW	Abstract	Point	Polygon	Individual	Category
Caduff and Timpf (2005)										■	■		■	
(Elias and Sester 2006)								■			■		■	
Dale et al. (2005)								■			■		░	░
Duckham et al. (2010)		■		■	░		■		■		■	░		■
Klippel and Winter (2005)										■	■		■	
Richter (2007), Richter and Klippel (2007)										■		■	■	

the given route. Consequently, approaches to landmark integration work on a functional level only (except for Duckham et al. 2010, discussed below); they usually take a set of landmark candidates to be given.

The separation between identification and integration also becomes apparent when looking at the chosen data sources. Each approach for landmark identification uses a concrete data source, mostly spatial databases. For landmark integration, many approaches are not specific regarding the kind of data source underlying their approach. Only Duckham et al. (2010) explicitly use a database of POIs in their case study that is taken from the *WhereiS* map server. Likewise, Dale et al. (2005) claim to base their approach on existing GIS data.

It can further be observed that with the exception of Duckham et al. (2010), all approaches employ individual features rather than categories. For the CORAL system (Dale et al. 2005), this is not really known, but it most likely uses individuals tagged manually based on common-sense assumptions about suitable categories (thus, the light gray marking in the table for this aspect). Finally, almost all approaches for identification and integration assume landmarks to function as point-like entities along a route, very much as defined by Lynch (1960). Duckham et al. (2010) acknowledged that other kinds of landmarks may be useful and presented some ideas on how to integrate them. Richter (Richter 2007; Richter and Klippel 2007) took this idea the furthest. He argued for the need to integrate linear and area-like objects in structuring route information (cf. Hansen et al. 2006) and developed a uniform approach to determining the functional role of landmarks with different geometries.

3 Challenges: Why are Landmarks Not Used in Commercial Systems?

Several challenges have prevented the integration of landmark references into commercial systems up till today. The calculation and generation of directions in these systems are based on simple, efficient algorithms. Metric distances and references to street names, as they are used today in commercial navigation software, are easily calculable from a geo-referenced network representation of the street layout. Landmarks need to be embedded into this existing network structure in a seamless way, i.e., the graph needs to be annotated with additional features such that they are easily integratable into the directions, ideally already in the path search. Some systems combine metric distances with references to traffic lights to provide additional context.[2] In some systems, points of interest (POIs), such as hotels or gas stations, are accessible. While these POIs could be used as sets of

[2] Often, however, this is done without taking into account the presence of other traffic lights. It is not uncommon to get instructions, such as 'in 500 m, at the traffic lights, turn left,' with another set or two of traffic lights before the one referred to.

landmark candidates in principle, they are hardly ever used for describing the route to take, but rather as commercial announcements or selectable destinations. Also, given the commercial nature of POIs, there will be a bias of employing specific types of landmarks only, such as fast food restaurants or gas stations, and their distribution and density will likely lead to great variations in the quality of landmark-based navigation services, as can be seen from the analysis below.

The seamless integration of landmarks requires a suitable data structure; the Urban Knowledge Data Structure (Hansen et al. 2006; Klippel et al. 2009) that is based on OGC's OpenLS specification[3] might be such an approach that can deal with different types of landmarks and offers mechanisms for structuring route information.

Given such a data structure, there is still the need to identify landmark candidates and then to integrate suitable candidates into route directions. However, as can be seen in Table 1, these two kinds of approaches often lack integration. Elaborate approaches to landmark integration either ignore the problem of identifying landmark candidates (Caduff and Timpf 2005; Richter 2007) or do not provide any details on how this is done (Dale et al. 2005). Some of the approaches to identification, namely Raubal and Winter (2002), provide some ideas on integrating landmarks into formal specifications of turning actions, but are restricted in the way references may be created and also fail to discuss situations where no landmarks are present.

Klippel and Winter (2005) explicitly combined identification and integration of landmarks by extending the Wayfinding Choreme grammar (Klippel 2003)—a formal specification of movement behavior in wayfinding—with landmark annotations. Consequently, this approach is listed both under identification and integration in Table 1. Winter et al. (2008) used the machine learning approach of Elias (2003) to identify landmark candidates, which are then used in generating destination descriptions (Tomko and Winter 2009). While these integrated processes work in theory, they are highly data intensive. They use individuals, i.e., identify individual features that may serve as a landmark. To gain useful results, these individuals need to be described in great detail, which is especially true for the calculation of façade salience in Raubal and Winter (2002). The required information is hard to collect automatically and, therefore, labor-intensive, will need to be specifically collected for each town and will result in large amounts of data. This makes it unlikely that it ever will appear in commercial databases due to the attached immense collection efforts and costs.

Therefore, looking at categories rather than individuals seems to be the more promising way, as can be seen with the implementation of Duckham et al.'s (2010) approach in the *WhereiS* web service. Using categories, properties of individual features do not need to be known since they are inferred by some heuristics from a general assessment of a specific category's suitability as landmark. Much less data is required; relevant information comprises location, geometry, and type of feature.

3 http://www.opengeospatial.org/standards/ols

Table 2 Example of Route Directions generated with whereis.com.au

	Distance	Time
Directions A		
Start: Bourke St, Melbourne, VIC 3000		
1. Continue on Royal La, Melbourne—head towards Bourke St at Red Violin	1 km	
2. Turn right onto Bourke St, Melbourne	0.6 km	1 min
3. Turn left onto Spring St, Melbourne *at Imperial Hotel*	0.1 km	14 s
4. Continue along Nicholson St, East Melbourne *at Princess Theatre @ Marriner Theatres*	0.7 km	1 min
5. Turn right onto Palmer St, Carlton *at Melbourne Museum*	14 m	2 s
6. Turn right onto Nicholson St, Fitzroy *at Academy Of Mary Immaculate Catholic*	0.1 km	8 s
7. Arrive at Nicholson St, Fitzroy		
Sub Total:	1.5 km	3 min
End: Melbourne Museum, 11 Nicholson St, Carlton, VIC 3053		
Total:	1.5 km	3 min
Directions B		
Start: Melville Rd, Brunswick West, VIC 3055		
1. Continue on Melville Rd, Brunswick West—head towards Bakers Pde	0.1 km	12 s
2. Turn right onto Moreland Rd, Brunswick West	1.6 km	3 min
3. Turn left onto Sydney Rd, Brunswick *at Moreland Hotel*	0.6 km	1 min
4. Turn right onto Rennie St, Coburg	0.8 km	1 min
5. Turn left onto Darlington Gr, Coburg	0.3 km	46 s
6. Turn right onto Carlisle St, Coburg	0.1 km	9 s
7. Arrive at Carlisle St, Coburg		
Sub Total:	3.5 km	6 min
End: Carlisle St, Coburg, VIC 3058		
Total:	3.5 km	6 min

Directions A from Bourke St to the Melbourne Museum, the route used as an example in Duckham et al. (2010); Directions B from Brunswick West to Coburg, two urban residential districts in Melbourne

Table 2 shows two sample route directions including landmark references that were generated using the *WhereiS* web service.[4] Directions A (from Bourke St to Melbourne Museum) contain several references to landmarks, which illustrates that this approach has great potential for commercial systems (note that the actual integration of landmarks, i.e., the generation of directions may still be improved for better conceptualization of the turning actions). However, as directions B (from Melville Rd to Carlisle St), which lead through two of Melbourne's urban residential districts close to the city center, illustrate, landmark candidates are not evenly distributed across the environment. The route described by directions B is comparable in length and complexity to the route of directions A, and they are not far apart from each other. Still, for route B, far less landmark candidates are available than for route A. Duckham et al. state that in the current *WhereiS*

[4] Accessed on March 29, 2010.

implementation for all of Australia only 170,000 features from 66 categories can be used as landmarks.[5]

This results in a sparse distribution of landmark candidates throughout the country; it can also be expected that the density is significantly higher in inner-city areas compared to suburbs or rural areas. Comparable results were found in a diploma thesis at the University of Bremen (Wagner 2009) that used a similar approach. An informal evaluation showed that all selected landmark candidates are sensible (i.e., landmarks are visible and identifiable along the route), but the geographic data set used contains too few features to properly cover large parts of an environment.

To sum up, while generating landmark references based on category information rather than on individual landmark properties seems to be the most prominent way to go in automatic landmark identification, a remaining challenge is to pool sufficient information about a sufficient number of features from a sufficient number of useful categories such that enough landmark candidates emerge to cover all parts of an environment.

4 Outlook: User-Generated Landmark Information

The more promising option to get at the missing landmark data is to tap into the vast and ever increasing repositories of user-generated content, a lot of which is geographic in nature (Goodchild 2007; Krumm et al. 2008; Sui 2008). User-generated content refers to data that is contributed to a service by its users. Usually, this data collection happens without a central authority managing or supervising the collection process. The individual approaches to data collection vary and cover a spectrum from conscious, dedicated user action ('volunteered') to rather passive modes ('citizens as sensors'). These approaches are made possible by the recent advent of new web technologies—commonly termed 'Web 2.0', or 'GeoWeb' in the spatial domain—and the ubiquity of network connectivity (see also Hirtle and Raubal, this volume).

These approaches to user-generated content can either be *indirectly* or *directly* exploited in landmark identification. Indirect approaches would tap into existing data sources of contributed data, similar to the web mining approaches discussed above, while direct approaches would create new sources specifically tailored to serve as sets of landmark candidates.

The GeoCAM project (Zhang et al. 2009), for example, aims at extracting meaningful parts from web documents containing route directions. These parts are the origin, the destination, and the instructions to get from one to the other. This

[5] To get a better idea of what this means: if landmark candidates were evenly distributed across Australia, there would be roughly 200 candidates within the area of Melbourne, or about 1 feature every 45 km^2.

extraction could be extended to also filter landmark information from these instructions, similar to Tezuka and Tanaka (2005), which is already hinted at in Zhang et al. (2009). Such landmark extraction from web documents would be an indirect approach. Others used geo-referenced (and/or) annotated photographs from photo sharing websites to identify landmarks; see also approaches to place identification from user-generated content that use similar mechanisms, e.g., Hollenstein and Purves (2010); Mummidi and Krumm (2008). For example, Schlieder and Matyas (2009) used the Panoramio photo database[6] to identify prominent sights in four European cities. Crandall et al. (2009) used Flickr[7] as a source to identify representative views of cities, which typically correspond to some salient geographic feature, such as the town hall or a cathedral. These identified sights are prominent, outstanding features in these cities, given that they have been photographed multiple times by different users. Thus, they can be expected to be salient, and may be used as landmark candidates. However, such an approach to landmark identification requires places actually being photographed or otherwise captured in user-generated content, which, again, most likely will lead to a sparse, uneven distribution of candidates. The major cathedral in the center of town may be photographed thousands of times, while the neighborhood churches in the suburbs may never appear in any photo collection. Thus, indirect approaches of exploiting user-generated content for landmark identification, while not suffering from the insurmountable costs of data collection (users essentially provide the data for free), still suffer from too few landmark candidates in large areas of an environment to be useful for navigation services.

Direct approaches to user-generated landmark content may result in a more even distribution of landmark candidates. For example, OpenStreetMap collects user-generated content to provide topographic data of the world. While accuracy and completeness is not the same everywhere, overall this project has managed in the last few years to create a data set that is comparable with authoritative data sets in at least the more densely populated areas of the Western world (Haklay 2010; Zielstra and Zipf 2010).

Further, several dedicated services providing spatial information for specific user groups have been suggested in the literature. Priedhorsky et al. (2007) proposed a Wiki-like service for bicyclists. Here, users directly contribute semantic information to enhance the bike riding experience. CityFlocks proposed to exploit detailed knowledge of people living in a neighborhood to annotate places of interest in that neighborhood (Bilandzic et al. 2008). The latter is not particularly geared towards navigation information, but rather to find and judge places to get food and other daily needs.

It is conceivable that applications similar to CityFlocks can provide landmark information (Richter and Winter 2011). They could tap into locals' knowledge and expertise to identify landmarks for navigation services. Such services would need

---

[6] http://www.panoramio.com

[7] http://www.flickr.com

to be designed such that users stay motivated to contribute (for a discussion of why people contribute, see e.g., Budhathoki et al. (2010) and would require carefully crafted instructions of what to do in order to avoid biasing users in what they contribute. Another direct way of identifying landmark candidates through user-generated content may be to ask users of a website to describe intersections seen on photographs of that intersection (e.g., taken from Google StreetView). This may be implemented either as a photo tagging game similar to Google Image Labeler[8] (Ahn et al. 2006) or as a (low-)paid job on websites for human intelligence tasks, such as Amazon Mechanical Turk.[9] Such approaches directly tap into humans' semantic knowledge of an environment. They mark some geographic features of a neighborhood as landmarks and, depending on how the user interface is designed, also collect reasons as of why these features are seen to be landmarks. However, as with most user-generated content, there is no guarantee that the provided information is actually useful. Thus, these approaches require the incorporation of trust and reputation mechanisms (Alfaro et al. 2011; Flanagin and Metzger 2008) in the creation of landmark candidate sets. They may also incorporate mechanisms to 'follow' landmarks of specific users (e.g., because they turn out to be especially effective for some users), this way enabling user-specific landmarks in navigation instructions.

In summary, research over the last decades has clearly established the important role landmarks play in our understanding of and communication about space. Empirical findings have inspired several computational approaches to the identification of landmark candidates and their integration into (automatically generated) route directions. However, these advances in basic research failed to find their way into commercial applications; landmarks are hardly ever considered in of-the-shelf navigation services. This failure can be attributed to two aspects: (1) the immense effort and, thus, costs attached to the acquisition of the required data for many of the approaches; (2) the highly skewed distribution of landmark candidates in available spatial data, which leaves large parts of an environment without suitable candidates. These challenges may be tackled by exploiting principles and methods of crowd-sourcing. In light of current developments in user-generated content, where users participate in building up and improving the (web) services they use, instead of investing in ever more complex computational approaches that rely on infeasible top-down, authoritative data collection methods, computational intelligence and smart interface design should be invested to achieve sustainable crowd-sourced landmark collection services that exploit human intelligence of (local) experts.

Acknowledgments A first draft of this chapter was written while the author was at the School of Information Science, The University of Pittsburgh. The ideas presented in the outlook greatly benefited from discussions at the School with Stephen Hirtle, Sherry Koshman, Hassan Karimi and Peter Brusilovsky, as well as later on with Stephan Winter at the University of Melbourne.

8 http://images.google.com/imagelabeler/

9 https://www.mturk.com/mturk/welcome

Martin Tomko provided further feedback on an earlier draft of this chapter. The author received support for attending the meeting in Las Navas from the Transregional Collaborative Research Center SFB/TR 8 Spatial Cognition which is funded by the DFG. Feedback by the participants of that meeting, especially Martin Raubal, is also gratefully acknowledged.

References

Ahn L, Liu R, Blum M (2006) Peekaboom: a game for locating objects in images. In: CHI '06: Proceedings of the SIGCHI conference on human factors in computing systems. ACM, New York, pp 55–64

Alfaro LD, Kulshreshtha A, Pye I, Adler BT (2011) Reputation systems for open collaboration. Commun ACM 54(8):81–87

Bilandzic M, Foth M, De Luca A (2008) CityFlocks: designing social navigation for urban mobile information systems. In: DIS '08: Proceedings of the 7th ACM conference on designing interactive systems. ACM, New York, pp 174–183

Budhathoki NR, Nedovic-Budic Z, Bruce B (2010) An interdisciplinary frame for understanding volunteered geographic information. Geomatica 64(1):11–26

Burnett G, Smith D, May A (2001) Supporting the navigation task: characteristics of 'good' landmarks. In: Hanson MA (ed) Contemporary ergonomics 2001. Taylor & Francis, London, pp 441–446

Caduff D, Timpf S (2005) The landmark spider: representing landmark knowledge for wayfinding tasks. In: Barkowsky T, Freksa C, Hegarty M, Lowe R (eds) Reasoning with mental and external diagrams: computational modeling and spatial assistance—papers from the 2005 AAAI spring symposium, Menlo Park, pp 30–35

Couclelis H, Golledge RG, Gale N, Tobler W (1987) Exploring the anchor-point hypothesis of spatial cognition. J Environ Psychol 7:99–122

Crandall DJ, Backstrom L, Huttenlocher D, Kleinberg J (2009) Mapping the world's photos. In: Proceedings of the 18th international conference on world wide web. ACM, Madrid, pp 761–770, 1526812

Dale R, Geldof S, Prost J-P (2005) Using natural language generation in automatic route description. J Res Pract Inf Technol 37(1):89–105

Denis M (1997) The description of routes: a cognitive approach to the production of spatial discourse. Cahiers Psychologie Cognitive 16(4):409–458

Duckham M, Winter S, Robinson M (2010) Including landmarks in routing instructions. J Locat Based Serv 4(1):28–52

Elias B (2003) Extracting landmarks with data mining methods. In: Kuhn W, Worboys M, Timpf S (eds) Spatial information theory. Lecture notes in computer science. Springer, Berlin, pp 375–389

Elias B, Paelke V, Kuhnt S (2005) Concepts for the cartographic visualization of landmarks. In: Gartner G (ed) Location based services & telecartography. Geowissenschaftliche Mitteilungen, TU Vienna, pp 1149–1155

Elias B, Sester M (2006) Incorporating landmarks with quality measures in routing procedures. In: Raubal M, Miller H, Frank A, Goodchild M (eds) Geographic information science. Lecture notes in computer science. Springer, Berlin, pp 65–80

Flanagin A, Metzger M (2008) The credibility of volunteered geographic information. GeoJournal 72(3):137–148

Goodchild M (2007) Citizens as sensors: the world of volunteered geography. GeoJournal 69(4):211–221

Haklay M (2010) How good is volunteered geographic information? A comparative study of OpenStreetMap and ordnance survey datasets. Environ Plan B Plan Des 37(4):682–703

Hansen S, Richter K-F, Klippel A (2006) Landmarks in OpenLS—a data structure for cognitive ergonomic route directions. In: Raubal M, Miller H, Frank AU, Goodchild MF (eds) Geographic information science. Lecture notes in computer science. Springer, Berlin, pp 128–144

Hirtle SC, Jonides J (1985) Evidence of hierarchies in cognitive maps. Mem Cogn 13(3):208–217

Hollenstein L, Purves R (2010) Exploring place through user-generated content: Using Flickr to describe city cores. J Spat Inf Sci 1:21–48

Klippel A (2003) Wayfinding choremes. In: Kuhn W, Worboys M, Timpf S (eds) Spatial information theory. Lecture notes in computer science. Springer, Berlin, pp 320–334

Klippel A, Hansen S, Richter K-F, Winter S (2009) Urban granularities—a data structure for cognitively ergonomic route directions. GeoInformatica 13(2):223–247

Klippel A, Richter K-F, Hansen S (2005) Structural salience as a landmark. In: Workshop mobile maps 2005, Salzburg

Klippel A, Winter S (2005) Structural salience of landmarks for route directions. In: Cohn AG, Mark D (eds) Spatial information theory. Lecture notes in computer science. Springer, Berlin, pp 347–362

Krumm J, Davies N, Narayanaswami C (2008) User-generated content. Pervasive Comput 7(4):10–11

Lovelace KL, Hegarty M, Montello DR (1999) Elements of good route directions in familiar and unfamiliar environments. In: Freksa C, Mark DM (eds) Spatial information theory. Lecture notes in computer science. Berlin, Springer, pp 65–82

Lynch K (1960) The image of the city. The MIT Press, Cambridge

May AJ, Ross T, Bayer SH, Tarkiainen MJ (2003) Pedestrian navigation aids: information requirements and design implications. Pers Ubiquit Comput 7(6):331–338

Michon P-E, Denis M (2001) When and why are visual landmarks used in giving directions? In: Montello DR (ed) Spatial information theory. Lecture notes in computer science. Springer, Berlin, pp 400–414

Mummidi L, Krumm J (2008) Discovering points of interest from users' map annotations. GeoJournal 72(3):215–227

Presson CC, Montello DR (1988) Points of reference in spatial cognition: stalking the elusive landmark. Br J Dev Psychol 6:378–381

Priedhorsky R, Jordan B, Terveen L (2007) How a personalized geowiki can help bicyclists share information more effectively. In: WikiSym '07: Proceedings of the 2007 international symposium on wikis. ACM, New York, pp 93–98

Raubal M, Winter S (2002) Enriching wayfinding instructions with local landmarks. In: Egenhofer M, Mark D (eds) Geographic information science. Lecture notes in computer science. Springer, Berlin, pp 243–259

Richter K-F (2007) A uniform handling of different landmark types in route directions. In: Winter S, Duckham M, Kulik L, Kuipers B (eds) Spatial information theory. Lecture notes in computer science. Springer, Berlin, pp 373–389

Richter KF (2008) Context-Specific Route Directions - Generation of Cognitively Motivated Wayfinding Instructions, DisKi 314 / SFB/TR 8 Monographs Volume 3. IOS Press, Amsterdam, The Netherlands

Richter K-F, Klippel A (2007) Before or after: prepositions in spatially constrained systems. In: Barkowsky T, Knauff M, Ligozat G, Montello DR (eds) Spatial cognition V. Lecture notes in artificial intelligence. Springer, Berlin, pp 453–469

Richter K-F, Winter S (2011) Citizens as database: conscious ubiquity in data collection. In: Pfoser D, Tao Y, Mouratidis K, Nascimento M, Mokbel M, Shekhar S, Huang Y (eds) Advances in spatial and temporal databases. Lecture notes in computer science. Springer, Berlin, pp 445–448

Ross T, Burnett G (2001) Evaluating the human-machine interface to vehicle navigation systems as an example of ubiquitous computing. Int J Hum Comput Stud 55:661–674

Ross T, May A, Thompson S (2004) The use of landmarks in pedestrian navigation instructions and the effects of context. In: Brewster S, Dunlop M (eds) Mobile computer interaction-MobileHCI 2004. Lecture notes in computer science. Springer, Berlin, pp 300–304

Sadeghian P, Kantardzic M (2008) The new generation of automatic landmark detection systems: challenges and guidelines. Spat Cogn Comput Interdisc J 8(3):252–287

Schlieder C, Matyas C (2009) Photographing a city: an analysis of place concepts based on spatial choices. Spat Cogn Comput Interdisc J 9(3):212–228

Siegel AW, White SH (1975) The development of spatial representations of large-scale environments. In: Reese HW (ed) Advances in child development and behaviour. Academic Press, New York, pp 9–55

Sorrows ME, Hirtle SC (1999) The nature of landmarks for real and electronic spaces. In: Freksa C, Mark DM (eds) Spatial information theory. Lecture notes in computer science. Springer, Berlin, pp 37–50

Sui DZ (2008) The wikification of GIS and its consequences: or Angelina Jolie's new tattoo and the future of GIS. Comput Environ Urban Syst 32(1):1–5

Tezuka T, Tanaka K (2005) Landmark extraction: a web mining approach. In: Cohn AG, Mark DM (eds) Spatial information theory. Lecture notes in computer science. Springer, Berlin, pp 379–396

Tomko M (2004) Case study—assessing spatial distribution of web resources for navigation services. In: 4th international workshop on web and wireless geographical information systems W2GIS04, Koyang

Tomko M, Winter S (2009) Pragmatic construction of destination descriptions for urban environments. Spat Cogn Comput Interdisc J 9(1):1–29

Wagner DB (2009) Bestimmung von Landmarken aus räumlichen Daten. Diploma thesis, Universität Bremen

Williams S (1998) Generating pitch accents in a concept-to-speech system using a knowledge base. In: Proceedings of the 5th international conference on spoken language processing, Sydney

Winter S (2003) Route adaptive selection of salient features. In: Kuhn W, Worboys M, Timpf S (eds) Spatial information theory. Lecture notes in computer science. Springer, Berlin, pp 349–361

Winter S, Raubal M, Nothegger C (2005) Focalizing measures of salience for wayfinding. In: Meng L, Zipf A, Reichenbacher T (eds) Map-based mobile services—theories, methods and implementations. Springer, Berlin, pp 127–142

Winter S, Tomko M, Elias B, Sester M (2008) Landmark hierarchies in context. Environ Plan B Plan Des 35(3):381–398

Zhang X, Mitra P, Xu S, Jaiswal AR, Klippel A, Maceachren A (2009) Extracting route directions from web pages. In: Twelth international workshop on web and databases (WebDB 2009), Providence

Zielstra D, Zipf A (2010) A comparative study of proprietary geodata and volunteered geographic information for Germany. In: Painho M, Santos MY, Pundt H (eds) Geospatial thinking: Proceedings of 13th AGILE international conference on geographic information science 2010