Automatic Generation and Application of Landmarks in Navigation Data Sets

Birgit Elias and Claus Brenner

Institute of Cartography and Geoinformatics, University of Hannover,
Appelstr. 9a, 30167 Hannover, Germany
birgit.elias@ikg.uni-hannover.de, claus.brenner@ikg.uni-hannover.de

Abstract

Landmark-based navigation is the most natural concept for humans to navigate themselves through their environment. It is desirable to incorporate this concept into car and personal navigation systems, which are nowadays based on distance and turn instructions. In this paper, an approach to identify landmarks automatically using existing GIS databases is introduced. By means of data mining methods, building information of the digital cadastral map of Germany is analyzed in order to identify landmarks. In a second step, a digital surface model obtained by laser scanning is used to assess the visibility of landmarks for a given route.

1 Introduction and Related Work

With the growing market for small mobile devices (PDA, mobile phones) the need for useful applications increases. One very important application is a guiding component for personal use. But also existing car navigation applications need improvement to make them more adapted to humans and thus more easy and reliable to use. Today's navigation systems give driving assistance in terms of *instructions* and *distances*, based on the current position and the underlying digital map, see Section 1.1. In contrast, research in the field of spatial cognition has shown that the use of landmarks is the most natural concept for humans to navigate themselves through unfamiliar environment. Therefore, it is very important to provide landmarks for navigation purposes to simplify the system for the user by giving more natural instructions. The concept of landmarks is explained in Section 1.2. For the automatic determination of landmarks different existing GIS databases are used, see Section 1.3.

In this paper, we will outline an approach to determine landmarks automatically in GIS databases. For that purpose, we subdivide the generation of landmarks in two different stages (see Section 2). First, we investigate the data for so called potential landmarks using data mining techniques (Section 2.1). After that, we narrow the selection down to route-specific landmarks considering their visibility from the point of view. Therefore, we need 3D data of the environment. Here, we use a DSM (Digital Surface Model) obtained from laser scanning data (Section 2.2). In Section 3 we discuss the results and give an outlook about future work.

1.1 Today's Car Navigation Systems

Modern car navigation systems have been introduced in 1995 in upper class cars and are now available for practically any model. They are relatively complex and mature systems able to provide route guidance in form of digital maps, driving direction pictograms, and spoken language driving instructions (Zhao 1997). Looking back to the first beginnings in the early 1980s, many nontrivial problems have been solved such as absolute positioning, provision of huge navigable maps, fast routing and reliable route guidance.

The maps used by car navigation systems not only contain the geometry and connectivity of the road network but also a huge amount of additional information on objects, attributes and relationships. A good overview can be obtained from the European standard *GDF*, see e.g. (*Geographic Data Files 3.0* 1995). Of particular interest are points of interest (POI) which include museums, theaters, cultural centers, city halls, etc.

Map data is acquired by map database vendors such as Tele Atlas or NavTeq and supplied to car navigation manufacturers in an exchange format (such as GDF). There, it is converted to the proprietary formats finally found on the map CD or DVD. This conversion is nontrivial since the data has to be transformed from a descriptive form into a specialized form supporting efficient queries by the car navigation system. Often, structures and values are pre-computed by this conversion process in order to relieve the navigation system's online resources such as bandwidth and CPU time.

1.2 Basic Theory of Navigation with Landmarks

The term landmark stands for a salient object in the environment that aids the user in navigating and understanding space. In general, an indicator of landmarks can be particular visual characteristic, unique purpose or meaning, or central or prominent location. They can be divided into three categories: visual, cognitive and structural landmarks. The more of these categories apply for the particular object, the more it qualifies as a landmark (Sorrows & Hirtle 1999). A study of Lovelace, Hegarty & Montello (1999) includes an exploration of the kinds and locations of landmarks used in route directions. It can be distinguished between four groups: choice point landmarks (at decision points), potential choice point landmarks (at traversing intersections), on-route landmarks (along a path with no choice) and off-route landmarks (distant but visible from the route). A major outcome of the study is that decision point and on-route landmarks are the most used ones in route directions of unfamiliar environments.

The use of landmarks in street maps is discussed by Deakin (1996). The findings reveal that when supplemental landmarks are given, the navigation process is more successful and less errors occur. Landmark symbolization represented by geometric symbols or stereotype sketches was found to be equally effective. Even in car navigation systems using graphic and voice instructions instead of maps, landmarks added to directional instructions were helpful. The determination of which kind of landmark is useful for a specific navigation task is investigated for the purposes of car navigation systems by Burnett (1998): especially 'road infrastructure', such as traffic lights and petrol stations are considered to be important. Prospective attributes of such landmarks include permanence, visibility, location in relation to decision point, uniqueness and brevity (Burnett, Smith & May 2001).

The automatic generation of landmarks is a matter of ongoing research. One approach uses the landmark concept of Sorrows & Hirtle (1999) to provide measures to specify formally the landmark saliency of buildings: the strength or attractiveness of landmarks is determined by the components visual attraction (e.g. consisting of façade area, shape, color, visibility), semantic attraction (cultural and historical importance, explicit marks such as shop signs) and structural attraction (nodes, boundaries, regions). The combination of the property values leads to a numerical estimation of the landmarks' saliency (Raubal & Winter 2002, Nothegger 2003). The concept was extended by a measure of advance visibility (Winter 2003). The visibility of façades while approaching to a destination point is determined. Combined with the measure of saliency it represents an approach to identify route-specific landmarks.

1.3 Used GIS Databases

In our approach we use the digital cadastral map of Lower Saxony in Germany. This database is an object oriented vector database of state-wide availability. The map includes information about parcels, land use, building and further administrative statements. We focus especially on building polygons and use geometry as well as additional semantic attributes like building use (residential, public use) and building labels (name or function of building such as church, kindergarten). For providing the road network, we used ATKIS data. Finally, the digital surface model (DSM) for this study was obtained by airborne laser scanning using the TopoSys scanner (Lohr 1999). The original point cloud was interpolated to a 1m $\times$ 1m grid. Last pulse data was used in order to reduce the influence of vegetation. The DSM is used for computing the visibility of objects, as outlined below in Section 2.2.

2 Determination of Landmarks

The identification of appropriate landmarks consists of two stages (see Figure 1): First, the detection of all potential landmarks in the geo-database. That means, the existing GIS database is searched for objects following our definition for a landmark: landmarks are assumed to be topographic objects which exhibit distinct and unique properties with respect to their local neighborhood. These properties determine the saliency of the objects, which in turn depends on different factors like size, height, color, time of the day or year, familiarity with the situation, direction of route, etc. For the first stage, the general geometric and semantic characteristics of the investigated objects in a certain neighborhood are needed only. This computation step is independent from the particular route chosen (see Section 2.1).

In a second step, the tentative selection has to be adapted to the route-specific needs such as visibility of the object from the decision point and the visibility while approaching this point. Finally, the position of the landmark relative to the route has to be considered in order to derive the appropriate route instruction which can be integrated into navigation data sets (see Section 2.2).

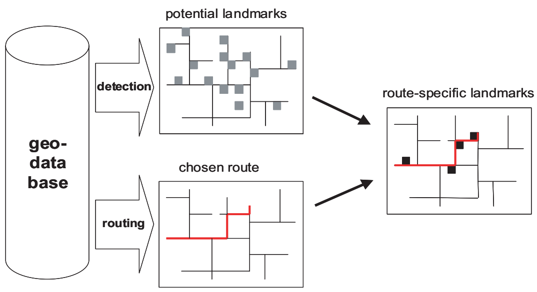

Fig. 1. Generation of landmarks.

2.1 Discovering Potential Landmarks

To make an automatic analysis process possible, we use data mining techniques to detect landmarks. Data mining methods are algorithms designed to analyze data, to cluster data into specific categories or identify regular patterns (Fayyad, Piatetsky-Shapiro, Smyth & Uthurusamy 1996). Basic models of data mining are clustering, regression models, classification, summarization, link and sequence analysis. These procedures can be applied to data sets consisting of collected attribute values and relations for objects. In order to identify landmarks among buildings, all existing information about the buildings has to be extracted: information about semantics (use, function) and geometry of the object itself (area, form, edges), but also information about topology, e.g. neighborhood relations to other buildings and other object groups (roads, parcel boundary etc.) and orientation of the buildings (towards north, next road, neighbor) are collected in an attribute-value table. Because of the findings of Lovelace et al. (1999) (see section 1.2), we only investigate (potential) decision points on the route for landmarks, that means all junctions in the underlying road network. For each potential decision point the local environment for the investigation is determined by means of a 360 degree visibility analysis to determine which objects are visible from that point of view at all. All selected buildings (creating the local environment) are transferred to the data mining process to detect the object with distinct and unique properties with respect to all others. We used the classification algorithm ID3 (Quinlan 1986) and the clustering approach COBWEB (Witten & Eibe 1999), see (Elias 2003). Here, we present the results of the first test using a modified ID3 algorithm.

Chosen Algorithm for Data Mining

Originally, the ID3 algorithm is a method of supervised learning for inducing classification models, also called decision trees, from training data. The decision tree is built from a fixed set of examples, each of which has several attributes and belongs to a class (like yes or no). The resulting tree can be used to classify future samples. The decision nodes of the tree are chosen by use of an entropy-based measure known as *information gain* (Han & Kamber 2001). Applied to our data the algorithm needs classified examples (the information which instances belong to the class 'landmark = yes'). The resulting tree provides the shortest optimal description possible for a classification into the given classes. As in our study we do not know the landmarks in beforehand, there are no classified examples. Thus, we use this procedure in a modified way: we enhance our data set with the class landmark (values: yes, no). Then, we iteratively hypothesize each building being a landmark, whereas all the other buildings are no landmarks. Afterwards, we process all these data sets with the algorithm and compare the resulting decision trees with each other. The assumption is that true landmarks are identified by yielding the most simple (shortest) description, which means a position in the decision tree close to the root (for further details see (Elias 2003)). At the moment, we only consider instances with a resulting decision tree with one branch as a potential landmark.

Preprocessing for Data Mining

Before processing the data with the data mining algorithm, a few preprocessing steps are necessary: first, we have to provide the attribute values for all buildings in the analysis environment (see (Elias 2003)). The used attribute

Table 1. Types of Attributes.

Nominal attributes	**Numeric attributes**
building use (residential, public, ...)	number of corners
deviation of corners being rectangular (yes, no)	building area
building functions (code table)	length, width of building and their ratio
orientation to street (along, across, angular)	orientation to north (angle)
orientation to neighbor (identical, angular)	distance to road
other parcel use than neighbors (yes, no)	number of neighbor buildings with direct contact
	ratio of building to parcel area
	built-up density around building
	number of buildings per parcel

values divide into two different types of attributes: nominal (values are categories) and numeric (see Table 1).

The derived attribute values have to be adapted to the chosen algorithm, because depending on the underlying mathematical concept only particular attribute types can be used. For example, the ID3 algorithm works with nominal attributes only. Therefore, a transformation between the attribute types is necessary. We have used the data mining package WEKA (Witten & Eibe 1999) which offers an automatic conversion from numeric to nominal values (by dividing the numerical values in different class ranges).

Processing of Potential Landmarks

We focus on the extraction of landmarks at decision points, that means junctions that can be potential turning points in the later routing description. Therefore, we use the road network of the ATKIS database to provide all possible junctions in the neighborhood. For each decision point we determine its local environment to investigate if there is a topographic object (in this approach, buildings) with distinct unique properties with respect to its local neighborhood, fulfilling our requirements for being a landmark. We use the results of the visibility analysis (section 2.2) to determine which objects are in fact visible from the point of view at the decision point (see Figure 2, left). The selected objects are used with their attribute values in the data mining process.

As a result, all instances which lead to a positive classification to the class 'landmark = yes' in the first branch are proposed as potential landmarks. In the example of Figure 2, right, the result of the modified ID3 processing is shown: Three different buildings are highlighted as being potential landmarks. They are chosen because each of them discriminates from the other buildings in one single attribute value: building number 1 has a different parcel use than its neighbor (because its a high voltage transformer building next to public buildings and a parking place). Object number 2 and 3 are singular because

Fig. 2. Decision point (black triangle). Left: selection of all visible buildings. Right: three potential landmarks after processing.

of their function: one is the cafeteria of the university and the second is a day-care center for children. The other large buildings near by the decision point could have been expected being potential landmarks, but they are all classified as university buildings in the data set and therefore they are not singular in their environment.

In Table 2 the degree of visibility (for a definition see Section 2.2), their distance to the decision point, and a short description about the building is given. It is clearly visible that object number 2 is better suited for being a landmark because of its high degree of visibility and nearness to the decision point. This has to be considered in a further assessment of the results.

Table 2. Degree of visibility.

No.	Visibility	Distance [m]	Description
1	4	216	high voltage transformer building
2	1278	60	cafeteria of university
3	3	118	kindergarten

In Figure 3 a panoramic camera view from the decision point is shown. Only potential landmark number 2, the University cafeteria, can be clearly recognized, because it has a high degree of visibility (the building is cut in two pieces, visible on the left and right margin).

2.2 Route-Specific Landmarks

Visibility Analysis

Landmarks can only be of use for navigation purposes if they are sufficiently visible during the actual navigation process. Although some conclusions on

Fig. 3. Panoramic camera view from decision point

visibility can be drawn from two-dimensional maps, important situations cannot be handled adequately. For example, Figure 4(a) shows this case where the visibility of the tower on the right is not revealed, as the height is not taken into consideration.
The optimal case is of course when a full three-dimensional city model is available. Then, visibility can be computed exactly, yielding even information on the visibility of single building faces. Buildings standing out behind other buildings are correctly identified (Figure 4(b)). Unfortunately, such three-dimensional city models are not always available, or only at substantial costs, which makes this approach not practical, especially with respect to the large areas typically covered in navigation databases.

However, we can do better if we base the visibility analysis directly on the DSM from laser scanning. We will not obtain "beautiful" visualizations but instead a sufficiently good estimate on which buildings can be seen from any given viewpoint (Figure 4(c)). We realized this approach as follows. For a given viewpoint, the position and viewing direction define the exterior orientation of a virtual camera of given horizontal and vertical viewing angle. This virtual camera represents the driver's view. The height is derived from the DSM itself, whereas the viewing angle can be obtained from the orientation of the corresponding street segment in the GDF or ATKIS dataset.

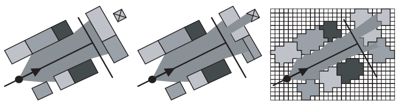

(a) (b) (c)

Fig. 4. Visibility analysis for buildings (gray boxes) standing along a street (black lines). The visibility cone is shown in dark gray. (a) Based on 2D ground plans. (b) Based on true 3D geometry. (c) Based on a discrete DSM.

The virtual image plane is then divided into a regular raster, each picture element (pixel) defining a ray in object space. All the rays are traced in object space to determine intersections with the DSM. For each hit, the corresponding object number is obtained by a lookup in an image containing ground plan id's. Figure 5 shows an example. It represents the same decision point as Figure 3. The objects appear in different shades of gray, which have been assigned randomly to the ground plan id's.

Fig. 5. Visibility computation for the scene shown above. Left: top view, showing the DSM and the cone of visibility. Right: virtual panoramic view where each shade of gray corresponds to an object (building) number.

For car navigation, a virtual camera should be used resembling a traditional (perspective) camera, as the driver's view is indeed constrained by the orientation of the car and the extends of the windshield. However, for personal navigation, the panoramic view shown here might be more appropriate, since the person to be guided might turn and look in all directions quite easily.

For each viewpoint, a list of records is obtained showing which objects are visible in the scene and how many pixels of the virtual image plane they cover. We use this number of pixels as a direct measure for the degree of visibility. Since each pixel corresponds to a well-defined area in the field of view of the person to be guided, the number of pixels is proportional to the absolute area in the person's field of view. Since the objects in the virtual view are directly linked to the id's of the map objects, one can immediately infer the degree of visibility for all objects in the cadastral map.

Visibility Tracking

In the last section, visibility was computed for a single view. However, landmarks selected for a routing instruction should be visible during an entire manoeuvre. This can be checked by tracking the visibility of objects along the trajectory defined by the corresponding manoeuvre. To this end, our algorithm traces the entire trajectory, generating virtual views at equidistantly spaced positions (2 meters distance in this case) and in the orientation defined by the trajectory. For each such view, the area covered by each object on the virtual image plane is determined.

Figure 6 shows a plot of all those areas along a trajectory passing the university cafeteria. One can see typical 'peaked' curves that are generated as objects appear, grow larger and finally disappear as the viewing position passes by. Those curves can be used to find out if an object is visible early

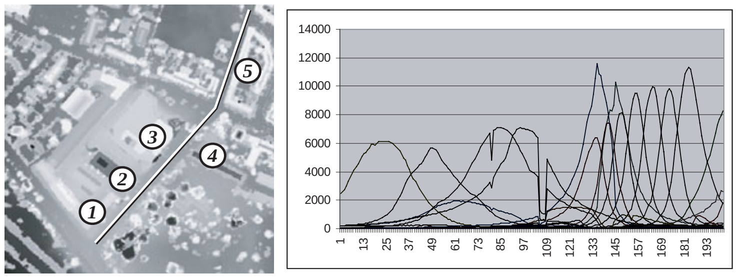

Fig. 6. Left: Location of a virtual trajectory (white line) passing the university cafeteria (marked with '3'). Right: Plot of all pixel counts along this trajectory (each abscissa step corresponds to 2 m real world distance).

enough, if it is visible throughout the entire manoeuvre, and if it covers enough area in the virtual image plane.

For example, in Figure 6, the first four peaks in the plot on the right correspond to the buildings marked with '1' to '4' in the top view on the left. Also, one can clearly distinguish the few, isolated peaks in the left half of the plot from the many, very dense peaks in the right half, caused by the second half of the trajectory which passes along a narrow street with many buildings standing closely together (marked '5' in the top view).

Integration into Navigation Databases

If an object is identified as being a landmark and additionally fulfills all requirements regarding its visibility, it can be used in a driving instruction. It is important to note that all the required computations can be done beforehand and there is no need to do them online on the navigation system itself – i.e., no cadastral maps or digital surface models are required in the end device. This can be accomplished by automatically investigating all junctions and all possible manoeuvres associated with those junctions, i.e. 'turn right', 'keep straight', etc. For each of those manoeuvres, the suitability of landmarks can be assessed.

In order to integrate landmark-based instructions into navigation systems, one has to simply extend the navigation matrices present in today's databases. Nowadays, those matrices identify which manoeuvres are allowed, which is done using a matrix of Boolean values. However, if we replace the entries by command codes which represent instructions such as 'turn right after the church', a landmark-based guidance can be derived. It is worth to note that no structural change in the database is necessary to achieve this.

3 Conclusion and Future Work

In the paper, we have shown a concept for the automatic extraction of landmarks from GIS databases. The current approach to determine potential landmarks with a modified ID3 algorithm has some drawbacks: even though singular objects are identified by the process effectively, problems arise if there are several similar objects (for example two churches) and each is qualified being a potential landmark at one decision point. The algorithm detects only objects which exist just once in the analyzed neighborhood, so double objects will never be processed as potential landmarks.

This is due to the fact that at the moment, the comparison of the decision trees is done on a very simple basis: we only consider the first branching in the decision tree. Therefore, a more complex analysis of the result trees has to be taken into account to avoid the disadvantage of deleting similar objects from the list of potential landmarks. Furthermore, it has to be evaluated if the used attributes are sensible and are properly weighted for the analysis process. For example, the function of the building has a very dominant impact on the results because of the nominal code list of function types. Furthermore, the use of the visibility information has to be improved and fully integrated: it is not sufficient to provide a generic visibility of objects, also the extent and distance of the object towards the decision point have to be taken into account. In the near future, we will also investigate into unsupervised clustering approaches and compare the results to the results obtained from using ID3.

Regarding the analysis of visibility, there are many possibilities for improvement. So far, we used the "virtual image size" to rate an object's visibility. However, from the virtual image, one can also obtain information on the distance, if the object is sticking out behind another, closer object, if it is part of the silhouette, and if it is closer to the center of view. First pulse laser scan measurements could be integrated to get a better approximation for the occlusion caused by trees. The DSM could also be used to feed additional information to the extraction, for example, small towers sticking out behind a larger building could be identified. The implementation of the visibility tracking could also use equidistant time sampling instead of space sampling, based on assumed vehicle speeds in the vicinity of intersections.

References

Burnett, G. E. [1998], Turn Right at the King's Head – Drivers' requirements for route guidance information, PhD thesis, Loughborough University.

Burnett, G., Smith, D. & May, A. [2001], Supporting the Navigation Task: Characteristics of 'Good' Landmarks, *in:* M. A. Hanson, ed., 'Contempory Ergonomics 2001', Taylor and Francis, pp. 441–446.

Deakin, A. K. [1996], 'Landmarks as Navigational Aids on Street Maps', *Cartography and Geographic Information Systems.*

Elias, B. [2003], Extracting Landmarks with Data Mining Methods, *in:* M. W. Werner Kuhn & S. Timpf, eds, 'Spatial Information Theory: Foundations of Geographic Information Science', Springer Verlag, pp. 398–412.

Fayyad, U. M., Piatetsky-Shapiro, G., Smyth, P. & Uthurusamy, R., eds [1996], *Advances in Knowledge Discovery and Data Mining*, AAAI Press/The MIT Press, Menlo Park, Californien.

Geographic Data Files 3.0 [1995], Technical report, European Committee for Standardization, CEN TC 278.

Han, J. & Kamber, M., eds [2001], *Data Mining: Concepts and Techniques*, Morgan Kaufmann.

Lohr, U. [1999], High Resolution Laserscanning, not only for 3D-City Models, *in:* D. Fritsch & R. Spiller, eds, 'Photogrammetric Week 99', Wichmann Verlag, pp. 133–138.

Lovelace, K., Hegarty, M. & Montello, D. [1999], Elements of Good Route Directions in Familiar and Unfamiliar Environments, *in:* C. Freksa & D. Mark, eds, 'Spatial Information Theory: Cognitive and Computational Foundations of Geographic Information Science', Springer Verlag, pp. 65–82.

Nothegger, C. [2003], Automatic Selection of Landmarks, Master's thesis, Technical University of Vienna.

Quinlan, J. R. [1986], 'Induction of Decision Trees', *Machine Learning*.

Raubal, M. & Winter, S. [2002], Enriching Wayfinding Instructions with Local Landmarks, *in:* M. Egenhofer & D. Mark, eds, 'Geographic Information Science', Vol. 2478 of *Lecture Notes in Computer Science*, Springer Verlag, pp. 243–259.

Sorrows, M. & Hirtle, S. [1999], The Nature of Landmarks for Real and Electronic Spaces, *in:* C. Freksa & D. Mark, eds, 'Spatial Information Theory: Cognitive and Computational Foundations of Geographic Information Science', Springer Verlag, pp. 37–50.

Winter, S. [2003], Route Adaptive Selection of Salient Features, *in:* M. W. Werner Kuhn & S. Timpf, eds, 'Spatial Information Theory: Foundations of Geographic Information Science', Springer Verlag, pp. 320–334.

Witten, I. H. & Eibe, F. [1999], *Data Mining: Practical Machine Learning Tools and Techniques with Java Implementations*, Morgan Kaufmann, San Francisco.

Zhao, Y. [1997], *Vehicle Location and Navigation Systems*, Artech House, Inc. Boston, London.