Determination of Landmarks and Reliability Criteria for Landmarks

Birgit Elias

Institute of Cartography and Geoinformatics
University of Hannover
Birgit.Elias@ikg.uni-hannover.de

Keywords: data interpretation, routing, landmarks, data mining, reliability criteria

Abstract

The navigation task is a very demanding application for mobile users. The algorithms of present software solutions are based on the established methods of car navigation systems and thus exhibit some inherent disadvantages: findings in spatial cognition research have shown that human users need landmarks for an easy and successful wayfinding. Typically, however, an object is not a landmark per se, but can be one relative to its environment. Unfortunately, these objects are not part of the todays route guidance information systems at the moment. Therefore, it is an aim of research to make landmarks for routing instructions available. In this paper we focus on a method to automatically derive landmarks from existing spatial databases. Here a new approach is presented by investigate existing spatial databases and try to extract landmarks automatically by use of a knowledge discovery process and data mining methods. Furthermore, a reliability criterion for landmarks to evaluate its uniqueness in the route environment will be developed.

1 Introduction

The wayfinding process in unknown environments is a very difficult and challenging task. Consequently, the interest to solve this problem by automated computed routing is enormous. Route guidance systems for vehicles are an established technology that are available even in middle class cars today.
Automatically produced routing directions are based on the data and concepts of current car navigation systems: the user gets turning instructions and metric distance measures. This concept works quite well, because of the restriction of automobiles to the road network. Currently, to improve usability navigation databases are enriched with additional information. Locations of petrol stations, pharmacies, public buildings and restaurants - so called points of interest - are part of the database. Another recent development is to transform the presentation of the route in the car into 3D-portrayals with models of historic buildings have been merged into the database to satisfy the touristic interest of the user.
The increasing amount of small and mobile technologies leads to a new group of user: the pedestrian. This evokes new problems for the automatic, computer driven navigation: pedestrians are not tied to the road network.
It makes no sense to simply adopt the current concepts of car navigation to mobile application for pedestrians. The discipline of spatial cognition investigates the human wayfinding concepts. Findings reveal, that humans need salient objects in the environment - so called *landmarks* - to orient and navigate themselves through space. But these objects are not per se a landmark, because being a landmark is a relative property depending on the local environment. Besides all that the capability of humans to estimate metric distances correctly are very poor.

Many research approaches in the field of spatial cognition have been undertaken about the wayfinding process and theory of landmarks, but only a few deal with the practical procedure to extract landmarks automatically from existing databases.

In this approach it is investigated, whether the salient objects in the environment can be extracted from spatial databases with methods of spatial data mining.
However, there is another influencing factor for the quality of a landmark which is not considered until now: What happens, if there is a similar landmark object on the path section towards the expected one? How much influence will a mistaken landmark have on the wayfinding process of the user and will this even trigger a wrong decision? Which uniqueness criteria does a landmark has to fulfil with respect to the whole route? At the moment there are no criteria for the certainty or reliability of landmarks depending on the whole route. In this paper a theoretic approach to calculate such an influence factor will be presented.

1.1 Related Work

1.1.1 Basic Theory of Navigation with Landmarks

There are two different kinds of route directions to convey the navigational information to the user: either in terms of a description (verbal instructions) or by means of a depiction (route map). According to (Tversky & Lee 1999) the structure and semantic content of both is equal, they consist of landmarks, orientation and actions. Using landmarks is important, because they serve multiple purposes in wayfinding: they help to organize space, because they are reference points in the environment and they support the navigation by identifying choice points, where a navigational decision has to be made (Golledge 1999). Accordingly, the term landmark stands for a salient object in the environment that aids the user in navigating and understanding the space (Sorrows & Hirtle 1999). In general, an indicator of landmarks can be particular visual characteristic, unique purpose or meaning, or central or prominent location.
Furthermore landmarks 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).
This concept is used by (Raubal & Winter 2002) 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, colour, visibility), semantic attraction (cultural and historical importance, explicit marks (e.g. shop signs)) and structural attraction (nodes (important intersections), boundaries (parting elements like rail tracks or rivers), regions (building blocks)). The combination of the property values leads to a numerical estimation of the landmarks' saliency.
A study of (Lovelace, Hegarty & Montello 1999) includes an exploration of the kinds and locations of landmarks used in 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 choice point and on-route landmarks are the most used ones in route directions of unfamiliar environments.

1.1.2 Principles of Data Mining

Data mining methods are algorithms designed to analyse data or to extract from data patterns into specific categories (Fayyad, Piatetsky-Shapiro, Smyth & Uthurusamy 1996). Basic models of data mining are clustering, regression models, classification, summarization, link and sequence analysis. Among other things the used algorithms are taken from the fields machine lerning, neural networks, rough set and fuzzy set theory or statistics (Cios, Pedrycz & Swiniarsky 1998).

Theses procedures can be applied to data sets consisting of collected attribute values and relations for objects. Attribute types can be nominal (values are categories with no ranking), ordinal (values have meaningful ordering) and numeric (values can be continuous, discrete, interval). Potential problems to be solved in the preprocessing stage are either incomplete or missing data in the data set or different attribute types to be handled.

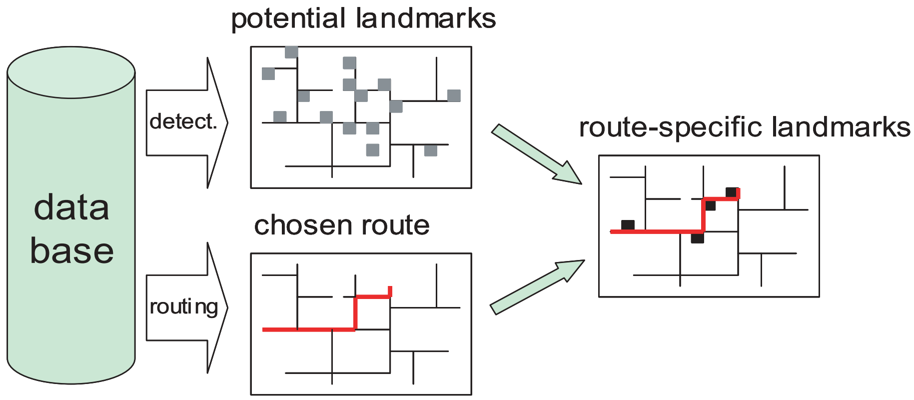

Figure 1: Determination of Landmarks

In the last years the scope of data mining from relational databases has been extended to spatial databases. The knowledge discovered through spatial data mining can be of various forms: providing characteristic and discriminant rules, extraction and description of prominent structures or clusters, spatial associations, and others (Koperski, Adhikary & Han 1996), (Anders & Sester 2000).

The algorithms can be devided in two basic techniques. According to the terminology of the machine learning community, there are methods for:

- learning from examples (*supervised learning*)
- learning from observation (*unsupervised learning*)

One very prominent example for supervised learning techniques is the ID3 algorithm by Quinlan (Quinlan 1986). It needs a set of classified examples and counter-examples and generates a decision tree, which can be used for classification of unknown examples. The method only deals with complete and nominal data sets and subdivides the training set of examples into homogeneous groups using discriminating attributes, whose selection is based on information gain calculation (Koperski et al. 1996). There are also extensions that enable to treat quantitative values as well. An application of knowledge acquisition for automatic interpretation of spatial data using ID3 is given in (Sester 2000).
The incremental conceptual clustering algorithm Cobweb is an example for unsupervised learning: clustering procedures divide examples in natural groups with similar features depending on a specific neighborhood concept. The chosen approach used here divides the examples into hierarchical groups. It is able to process with data records in form of symbolic attribute-value lists and builds up a decision tree (Witten & Eibe 1999).

1.2 Overview

The determination of landmarks includes two different phases: the detection of potential landmarks in the digital database and the exploitation of those that are relevant for a particular route (figure 1).
The detection of landmarks is completely independent of the chosen route. It depends only on the general geometric and semantic characteristics of the investigated objects and the defined neighborhood used for the analysis process. This computation can be done in pre-processing and provides all potential landmarks in the chosen environment.
In a second step, those landmarks are exploited that are relevant for the particular route. Now route-specific criteria are important, such as visibility, distance to route, particular orientation to route ("behind the church, turn right" and not "on the left side is the church, there turn right") and the uniqueness and reliability of the landmark in its neighborhood to avoid misleading.
In this paper we want to focus on some parts of the complete framework for landmarks determi-

nation: In section 2 we present an approach to identify landmarks automatically with data mining techniques to provide the potential landmarks. Secondly, in section 3 first investigations to develop reliabiltiy criteria to measure the quality and uniqueness of the landmark in its surrounding are introduced.
Methods to compute the visibility of objects according to a particular route are not part of this work. An approach using laserscanning data for analysing the visibility of landmarks is presented in (Brenner & Elias 2003).

2 An Approach of Identifying Landmarks

The above mentioned approach (section 1.1.1) to enrich wayfinding instructions with landmarks (Raubal & Winter 2002) follows the categories visual, semantic or structural attraction, introduced by (Sorrows & Hirtle 1999). Additionally, for grouping the structural elements the acknowledged structure elements of a city (nodes, boundaries, regions) named by (Lynch 1960) are used. The result is a catalogue of properties for each attraction element, which has to be filled with real data to make an automatic identification of landmarks possible. A problem could lie of course in data retrieval. Either there exists no database with the needed information or it is not available for greater areas. Maybe it is distributed over different databases and the collecting of all is too time-consuming and expensive like in the case of points of interest.
One way is to investigate existing topographic data sets: we investigate the information of available digital databases and extract the salient objects by analysing the real data. A first study to use existing databases for landmark extraction is given in (Elias 2002). In that study a set of specific feature types have been characterized as landmarks and have subsequently been extracted from the data sets. The problem is, however, that the quality as a landmark is relative: a multi-storing building may be an excellant landmark in a small city with only one such building - in New York however, it would be useless as a navigation aid. Therefore, in this paper the relative uniqueness of an object in its environment is taken into account. Thus, we start to investigate data mining methods for discovering landmarks using unsupervised learning methods.

The idea is that objects, which have a unique attribute in a certain environment, qualify as landmarks. Therefore, the underlying model is to compare the attribute values of all data records: These objects with distinct or even unique values differing from the global mean have to be something special. The procedure will also lead to an attribute ranking according to their importance for the model. If the chosen attributes are suitable for developing a global object schema and outliers from this schema present something particular, it is possible to determine landmarks through statistics and data mining methods.
As the goal is to extract the landmarks from a database, we need a data set. In this case we concentrate on the building objects. We enrich our explicit given information about buildings by deriving attributes and relations with the help of spatial analysis. The combination of different attribute-values leads to derived attributes, e.g. building length to width ratio. The data have to be preprocessed to unify the different attribute types, missing values have to be handled.
There are several open-source data mining and statistics tools available in the internet, for our investigation we use the WEKA implementation (Witten & Eibe 1999). This tool provides different data mining modules including preprocessing options.

2.1 Landmarks discovery in building data

2.1.1 Data Source and Selection

In this paper we focus on the digital cadastral map of lower saxony to use the benefits of an object oriented vector database of area-wide availability. This digital map includes buildings, parcels and their land use. Besides the geometry of the objects, the following attributes are available:

- building use: residential, public, underground, outbuildings
- land use types: public purposes, residential, commerce and service, industrial, mixed land uses, traffic, park, garden, sports, etc.

- building labels: name or function of building (e.g. town hall, kindergarten, church)
- special building parts with a roof: winter garden, car port etc.

To keep the first approach on a simple level, we simplify the conditions. We assume that for this analysis process the objects present a homogeneous group with predominantly identical attributes and that it is not possible to compare buildings with roads for example, because they have both important attributes not existing in the other group (building use: residential, official, etc.; road classification: highway, community road, etc.). In this first step, we decide to investigate the group of buildings und compose a complete attribute-value list for each building object.
The derived attribute values are calculated by combinating all explicit given information of the used data sets. In most cases the provided attributes describe geometric or topologic properties of the buildings, additionally some semantic information is given. Whether these describing characteristics of buildings are visually perceptible is not part of this study. It has to be tested separately, whether the calculated landmarks are perceived as real landmarks in reality.

The next step is to determine the situation, in which a landmark is needed for the wayfinding process and to convey the spatial cognition findings to our study. According to (Michon & Denis 2001) and (Lovelace et al. 1999) pedestrians simply progress along a route by direct themselves towards a landmark, thus landmarks are especially needed, when a turn on the route is required. Therefore, we try to provide a landmark at every (potential) choice point (or close to it) on the route. At every turning point of the route we start an examination, whether there is an applicable object in our data sets.
The amount of data makes it necessary to select the affected data sets with respect to the users

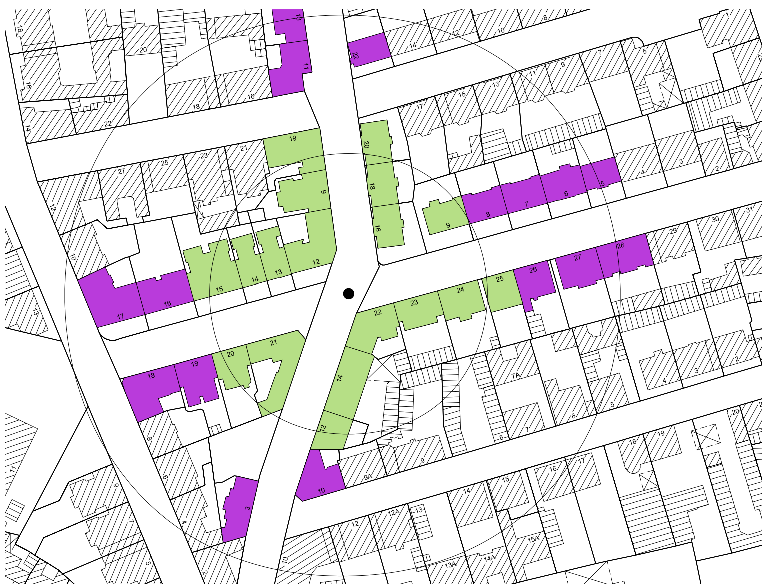

Figure 2: Selected buildings: neighborhood 50 m and 100 m

perspective and visibility within a predefined neighborhood. An overview is given in figure 2 showing a city situation with two different neighborhoods. The inner circle represents a distance of 50 m from the point of view and contains in this case nearly 20 buildings, the outer circle stands for a 100 m radius and more than 30 buildings, which can be used for searching for a salient one. Depending on the built-up area situation different kinds of visibility zones or neighborhoods can be established. If there is a dense, cropped situation a narrow neighborhood (for example the 50 m distance) is sufficient for the analysis.
In this szenario we disregard the fact that users usually move towards a certain direction and therefore want a landmark in front of him, not in his back. We only investigate the situation, that the user is standing in the middle of the road and searches for landmarks to support his current decision and guide his next movements.

2.1.2 Preprocessing: Providing Characteristics of Buildings

With respect to the formerly mentioned cadastral data set, the attribute-value table has to be composed for the selected buildings. All existing information about the objects has to be extracted from the database: information about semantics and geometry of the object itself, but also information about the topology, in this case neighborhood relations to other buildings and other object groups (for examples the distance to roads, to the parcel boundary etc.) are collected.
The derived attributes available from the digital cadastral map are shown in table 1. One profound disadvantage of the database content is the lack of height values for the buildings, which can be provided by laserscanning data (Brenner & Elias 2003)

no	graphic	attribute	description
1	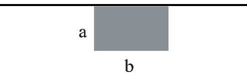	building area	maximum length b * maximum width a in $[m^2]$
2		building form	derivation to typical building form; ratio length/width
3	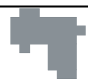	number of corners	counting quoins (typical: 4 or 6)
4	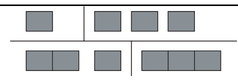	adjoined or detached building	detached, semi-detached, adjoined
5		distance to road	closest distance [m]
6	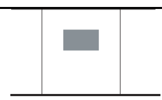	ratio of building area to parcel area	$\frac{building\ ground\ area}{parcel\ area}$
7	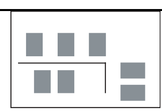	density of buildings (local neighborhood)	$\frac{number\ of\ buildings}{area\ (radius\ 100m)}$
8	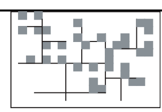	density of buildings (district)	$\frac{number\ of\ buildings}{area\ (radius\ 500m)}$
9	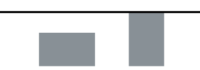	orientation to road	along (length towards road), across (width), angluar, building at corner [rad]
10	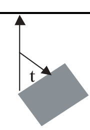	orientation to north	angle building length to north t [rad]
11	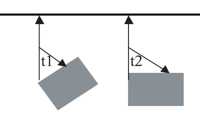	orientation to neighbor	difference angle to neighbour (t1 - t2) [rad]

12		perpendicular angle in building	derivation of angles to normal [rad]
13		number of buildings on parcel	counting buildings
14	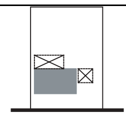	special building objects on parcel	counting number of car ports, winter garden etc.
15	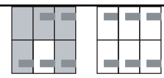	neighbor parcel land use	difference to surrounding neighbors: yes, no
16	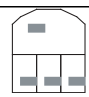	form of parcel area	number of corners, number of neighbors, adjoining roads

Table 1: Attributes and relations of buildings

2.1.3 Data Preprocessing

Table 1 provides values in different units: there is information in nominal categories (building use: public, residential, ...) or ordinal values (number of neighbors: 0, 1, 2) or even different metric values (distance to road [m], orientation to north [rad]). Depending on the chosen data mining algorithm different attribute types and units are needed, so the data has to be transformed. One procedure is for example the changeover from metric values to categories: the size of the buildings can be given in squaremetres or size classes can be set up (small, large). We have to take care of the transformation of different attribute types, because the type of transformation may influences the results of the whole analysis process. The specifications of needed attribute types change from algorithm to algorithm, but in data mining software automatic transformation operations are available (Witten & Eibe 1999).

To avoid the above mentioned problems and make the interpretation of the results easier, we simplify our approach in this first study by creation and use of a synthetic data set, which relies mostly on the existing data model of cadastral data. We build up a data set with limited attributes of the fictitious situation given in figure 3. As described in section 2.1.1, we assume that the user stands at a decision point and searches for a landmark in immediate environment. In this case the next neighborhood in a specified radius consists of 10 buildings used for the extraction procedure. The single objects are characterized by the following features with nominal values:

- building use: residential, public, outbuilding
- size of building: small, large
- number of immediate neighbors: 0, 1, 2
- orientation towards road: parallel, across, angular, corner
- distance from road: 0 or 3 metres
- building height: 12, 15 or 17 metres

Obviously, building number 3 is somehow singular and different to all other buildings: it has no direct contact to neighbored buildings and is the only large sized object. Thus, in the local environment this building is something particular and we expect the data mining methods to reveal this knowledge.

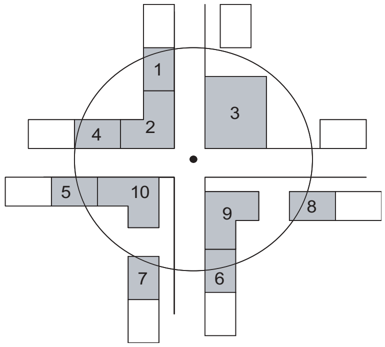

no	building use	size	neigh.	orientat.	dist.	height
1	residential	small	2	parallel	0	15
2	residential	small	2	corner	0	15
3	residential	large	0	corner	0	15
4	residential	small	2	parallel	0	15
5	residential	small	2	parallel	0	15
6	residential	small	2	parallel	0	15
7	residential	small	1	parallel	3	15
8	residential	small	1	parallel	3	15
9	residential	small	1	corner	3	15
10	residential	small	1	corner	3	15

Figure 3: Synthetic test data

2.1.4 Data Mining

We apply our test data set to the introduced data mining algorithm in section 1.1.2: a clustering approach using Cobweb. The used software WEKA is described in (Witten & Eibe 1999). Cobweb is a hierarchical clustering algorithm and thus an unsupervised learning method. Using this technique needs no explicit examples. Unclassified examples are parted in a hierarchy of natural groups by this procedure. The procedure processes the attribute values as nominal categories, thus no transformation is needed. Landmarks are objects that are in some sense unique in a given environment. As a clustering algorithm like Cobweb identifies clusters of objects with similar characteristics, an ideal landmark therefore should show up by being a singular object in one cluster.

The results of our data set computed with Cobweb are given in the dendrogram in figure 4. The depicted dendrogram shows the discrimination of the 10 instances in different branches. On the first level the records are diverted in three clusters consisting of 1, 4 and 5 instances. The latter two were further separated in two clusters containing only instances containing completely identical examples, where all attributes have the same values. The boxes enumerate the values of all attributes in the given cluster, below the boxes the corresponding instances are referenced to make a comparison with figure 3 possible. As it is clearly visible, the procedure aggregates similar

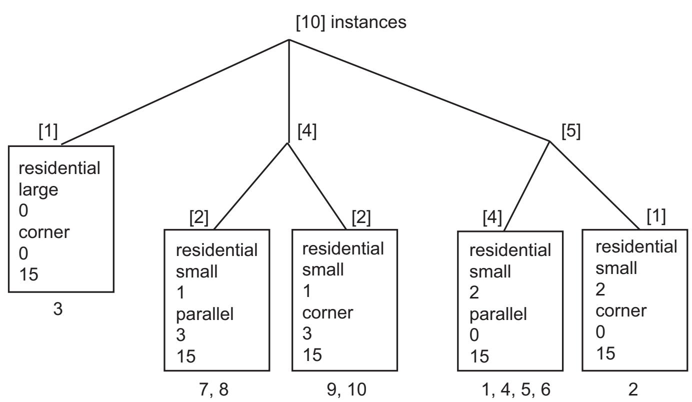

Figure 4: Results of COBWEB

instances into clusters. The more deviations in attribute values exist, the greater the distance in the hierarchical tree gets. Considering the cluster of the first level (see figure 5) the grey circles reveals the similiarity of the instances within the clusters. The groups built with 4 and 5 instances on the first level have 5 of 6 attribute values in common. Comparing the clusters of level one to each other, decreasing similiarity becomes clear. There are 2 or 3 attribute values identical, respectively this difference leads to the subdivision into 3 branches. The attributes of building number 3 differ very clearly from the other instances and therefore this record is isolated from the other clusters in one of the first levels.

Since the algorithm subdivides the records into similar groups, an instance with strongly different

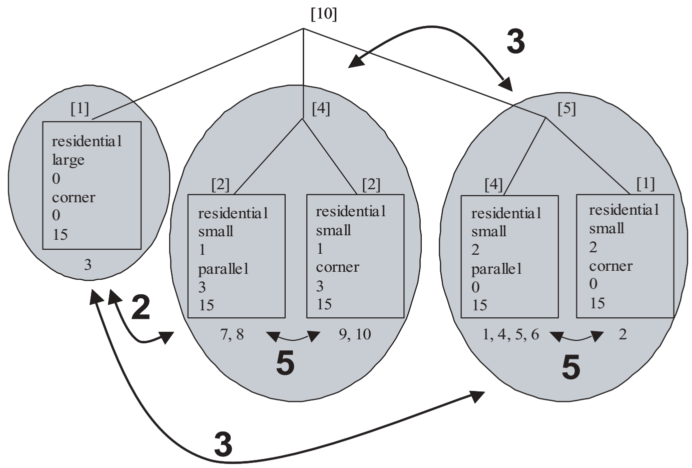

Figure 5: Coweb decision tree: attribute values in common

attribute values is separated from the others at a very high level in the decision tree. Because of its singularity, it is all alone in its group.

3 An Approach of Providing a Reliability Criterion

In this section we want to determine an quality measure for a potential landmark depending on the chosen route. This reliability criterion is based on considerations that a landmark is only useful if there is not another same looking object on the route misleading the user. Such a ”false” landmark is an object fitting to the description of the ”real” landmark (the object referred to in the route description) and lying on the followed route or near to it.

The investigation starts with the following initial assumptions of the given route direction:

- The route is part of a graph network (consisting of nodes and edges).
- The route description consists of a starting and end point, landmarks, as well as the turning instructions in terms of local orientations (left, right).
- Landmarks are only given at decision points, where the user has to make a turn.
- The moving direction between two landmarks is always relatively straight on until the next landmark is reached.

In this study on-route landmarks (between two decision points) are not used and the distance and position of the landmarks relative to the route are neglected (that means the impact of a ”false” landmark (that means a misleading object) situated on the wrong side is equal to one on the expected side).

3.1 Determination of Reliability Criteria

We assume a situation as portrayed in figure 6. A part of the routing graph is depicted consisting of the edges e_1, e_2 and e_3, connected through the nodes n_1, n_2. It starts at the last visited landmark (L_1) and leads to the next landmark (L_2), in this case a church. Because of the convention of going only straight on between two landmarks, the path segment is part of a ”stroke”. This term was coined for the road network generalization task by (Richardson 1999). It refers to a line consisting of several road segments put together to one feature, because the angles between two back-to-back segments are so small all segments are perceived by humans as one contiguous line. In this case, there are three strokes Str_1, Str_2 and Str_3.

Our goal is to determine a quality measure for the following landmark "church" depending on all other churches existing in a defined neighbourhood and its influence on the wayfinding task. To define the term neighbourhood we observe the moving behaviour of the user and its "errors": In case he follows the route instruction carefully, he uses only the edges and nodes that are part of the routing graph, that means he visits only stroke Str_1. This is our E_0-neighbourhood (E = error). If the user unconsciously makes one mistake and takes a turn at one of the nodes n_1, n_2 onto the edges e_4, e_6, e_8 or e_{10}, she enters another stroke (Str_2 or Str_3) and follows this one straight on expecting the show-up of the church. Only at the end of the stroke, when there is no possibility left to go further straight on and no expected landmark is in sight, the user notices his mistake. Thus the combination of the original routing graph plus all strokes branching at the original nodes builds the E_1-neighbourhood. The continuation of this consideration leads to the E_2 -neighbourhood and so on.

The hypothesis is that a "false" landmark showing up along the track misdirects the user, because it causes a wrong decision or confirms him being still on route misleadingly. Therefore all nodes in the defined neighbourhoods have to be checked if there exist "false" landmarks, that easily can be confused with the expected one.

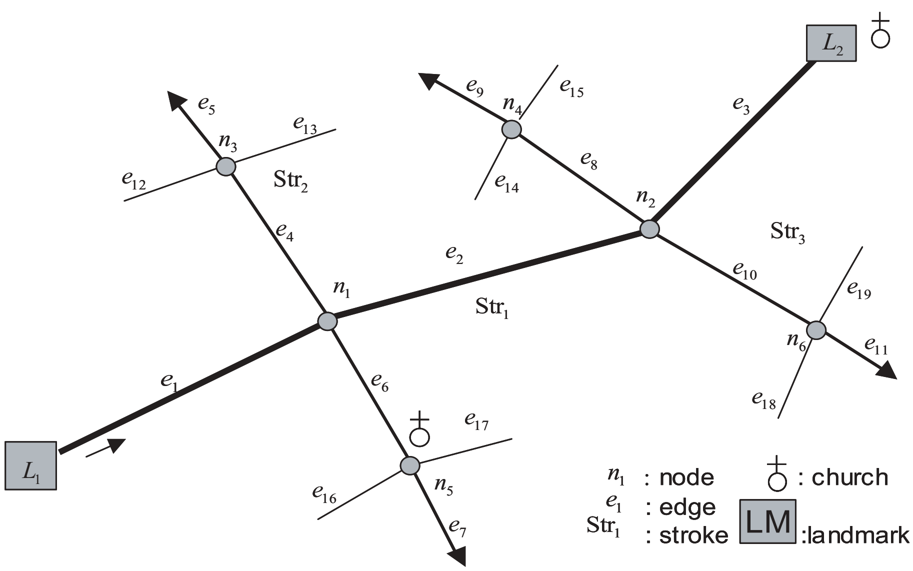

Figure 6: Routing graph between two landmarks and its first and second neighbourhood

It can be presumed that the impact of false landmarks in the E_0-neighbourhood, that means on the direct route segment to the next landmark, is much higher than the impact of a false landmark showing up in the E_1-neighbourhood. This has to be taken into account using weighting functions between the considered neighbourhoods. But also the absolute quantity of false landmarks, that means the number of equal objects in the chosen neighbourhood has an influence of the reliability of the correct landmark. The developed function results in a measure, that makes it possible to compare two or more different landmarks to select the best suited of them. If there is only one landmark, the critical value has to be determined, which trigger a more detailed description of the landmark (e.g. "church with two steeples") to avoid confusion with "false" landmarks. The developed function results in a measure that is able to take the uniqueness of a landmark into account - not only with respect to its immediate local neighbourhood, but also concerning the whole route.

4 Conclusion and Future Work

This paper presented an approach to automatically discover landmarks in spatial databases and first investigations to provide a quality measure for landmarks. In this study we concentrate on a building database. The clustering algorithm Cobweb is used to analyse synthetic test data of a spatial situation. The results reveal the potential of the method for discovering landmarks: they are characterized by beeing salient in their environment and the procedures reveal such singularity of objects by a cluster separated at a very high level containing only one element. The method seems to be suitable for the discovering process. To check the findings, whether they can be generalized for more applications, they have to be verified by applying it to real data basing on the proposed attributes in table 1. Another study has tested the synthetic data set with a modified ID3-algorithm and shows similar findings (Elias 2003).
The results with these data mining methods look very promising. Both lead to the desired result of identifying locally salient objects, that are likely to be distinguishable by (a set of) simple attributes. With Cobweb only the landmark object is identified. Its (visually) characterizing attributes have still to be extracted.
The data mining methods provide a mechanism to fully automatically extract objects with a relative uniqueness in a given environment. It is very flexible as it allows the introduction of arbitrary objects with arbitrary attributes - and still leads to the identification of the most unique object adapted to a given situation.

The investigations about quality measures for landmarks are on a very early stage. The reliability of the extracted landmarks has to be determined by a quality measure to avoid ambiguous landmarks misleading the user.
The long-term objective is to provide a framework to discover and assess the quality of landmarks within spatial databases in one single step.

References

Anders, K.-H. & Sester, M. [2000], Parameter-Free Cluster Detection in Spatial Databases and its Application to Typification, *in:* 'International Archives of Photogrammetry and Remote Sensing', Vol. 33(Part B4/1), ISPRS Congress, Amsterdam, pp. 75–82. Comm. IV.

Brenner, C. & Elias, B. [2003], Extracting Landmarks for Car Navigation Systems Using Existing GIS Databases and Laser Scanning, Technical report, ikg - Institute of Cartography and Geoinformatics, University of Hannover. (submitted to PIA 2003).

Cios, K., Pedrycz, W. & Swiniarsky, R. [1998], *Data Mining Methods for Knowledge Discovery*, Kluwer Academic Publishers.

Elias, B. [2002], Automatic Derivation of Location Maps, *in:* 'GeoSpatial Theory, Processing and Applications', Vol. 34/4, Ottawa, Canada.

Elias, B. [2003], Extracting Landmarks with Data Mining Methods, Technical report, ikg - Institute of Cartography and Geoinformatics, University of Hannover. (submitted to COSIT 2003).

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.

Golledge, R. D. [1999], *Wayfinding Behavior*, John Hopkins Press, chapter Human Wayfinding and Cognitive Maps, pp. 5–45.

Koperski, K., Adhikary, J. & Han, J. [1996], Spatial Data Mining: Progress and Challenges, *in:* 'Research Issues on Data Mining and Knowledge Discovery', Montreal, pp. 55–70.

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.

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

Michon, P.-E. & Denis, M. [2001], When and Why Are Visual Landmarks Used in Giving Directions?, *in:* D. Montello, ed., 'COSIT 2001', Springer Verlag, pp. 292–305.

Quinlan, J. R. [1986], 'Induction of Decision Trees', *Machine Learning* **1**, 81–106.

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.

Richardson, R. T. D. [1999], The "'good continuation"' principle of perceptual organization applied to the generalization of road networks, *in:* 'Proceedings of the ICA 19th International Cartographic Conference', Ottawa, Canada, pp. 1215–1223.

Sester, M. [2000], 'Knowledge acquisition for the automatic interpretation of spatial data', *International Journal of Geographical Information Science* **14**(1), 1–24.

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.

Tversky, B. & Lee, P. [1999], Pictorial and Verbal Tools for Conveying Routes, *in:* C. Freksa & D. Mark, eds, 'Spatial Information Theory: Cognitive and Computational Foundations of Geographic Information Science', Springer Verlag, pp. 51–64.

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