

Case Study - Assessing Spatial Distribution of Web Resources for Navigation Services

Martin Tomko¹

Department of Geomatics, University of Melbourne, Australia
CRC for Spatial Information
m.tomko@pgrad.unimelb.edu.au

Abstract. We present the results of a case study assessing the suitability of the World Wide Web in its current state of the art to provide supporting information for an automated pedestrian urban navigation service. Our experiment simulates an application providing a navigator with navigation aids—routemarks and landmarks—along a path generated by a web service. The instructions generated by this service are tested in human subject testing and evaluated. We conclude that automated, agent based systems could profit of the current state of Internet content to provide enhanced navigation services using landmarks. Further enhancements will be possible with the advent of the Semantic Web.

1 Introduction

Building specialised spatial databases for navigation is a tedious and lengthy process, leading to results requiring frequent updating, as the data soon become obsolete. However, the Internet provides almost whole world coverage with an accompanying unmatched wealth of data. We investigate the representation of the real world phenomena in the World Wide Web (Web), as opposed to the assessment of the Internet real estate geography. This case study illustrates how the spatial coverage and level of detail of the Web content can provide additional information to current pedestrian route planning web-services. It explores how wayfinding instructions from web-based urban navigation services can be enhanced by routemarks and landmarks on decision points, found on the web.

Our case study is a simulation of the actions that could be typically performed by an agent-based search engine automatically. Routemarks and landmarks were searched manually on the Web. Consecutively, routing instructions provided by the service were replaced by human generated instructions using landmarks and routemarks found, and their usability was assessed during wayfinding. This experiment provided insight in the capabilities of the Web and will lead to formalisation of the procedures and possibly an automated landmark-based navigation system.

This paper follows by introducing the problems of georeferencing web resources in Sect.2, follows by an introduction to human wayfinding with landmarks in Sect.3, continues with the description of our experiment, its environment and tools used in Sect.4. Experiment outcomes are presented in Sect.5. We conclude with a discussion on the results, problems and future work in Sect.6.

2 Assessing the Spatial Coverage of the Web Content

2.1 Spreading the Web

The Internet became an elementary part of everyday life for over half a billion people worldwide [1]. In Australia itself, over 10 million people, representing over 50% of the population, had access to the Internet in 2002. As published by Zook [2], Australia ranked 11th in the world in 2001 in proportion of online users to the world population (1.3%) with (23.78%) of registered domain names per capita. These numbers increase on a daily basis. Despite the “digital divide” between the more developed countries and the remaining of the world, the Internet is already easily accessible and contains information about nearly all aspects of life and of the world. Most of the websites are hosted on servers situated in one of the wealthier (OECD) countries, where the concentration of users is highest.

2.2 Web Content

The unmatched penetration of the Internet over the world and human lives is mirrored by the spread of the Web content. It is hard to think about any form of art, scientific endeavor, or private life that is not somehow represented on the web. It provides unparalleled freedom in terms of access, costs of content and use. This was reflected by unprecedented boom of “open” and “free” software, services and content (pieces of art, scientific results, news, databases, encyclopedias). This can hardly be matched by institutional building of expensive databases and softwares, with low flexibility of modification and updates, a problem shared by all creators of spatial databases.

Our experiment relies on the use of freely available content and services. This approach may, however, provide lesser reliability due to the uncertain quality of data used, as opposed to output generated by commercial, dedicated services and databases. The free content proved to be usable for our case study, and data more advanced filtering techniques can improve further the reliability of the output.

2.3 Mapping the Web

Several researches focused on the density of users of Internet in different countries [2, 3], the number of registered domains and individual hosts, and on the spatial distribution of addresses in the world or its regions [4]. These researches are valuable for assessment and visualisation of the spread of the Internet infrastructure [5] and its usage patterns. However, those efforts don't provide any information about the spatial coverage of real world phenomena through data available on the web. Detailed assessment of its spatial coverage is currently not feasible, until a sufficient proportion of the Web meets the Semantic Web [6], or better the Geospatial Semantic Web [7] capabilities. This will allow precise and machine readable coding of the resource location, and will provide a reliable mean to query its content, including geographic relation operations [8].

Though, it is estimated that 20-35% of all the searches performed on the web seek geographically relevant results [9]. The location of the host IP address is at best in indirect relation with the content served, and the localisation accuracy is unreliable [10]. This makes the retrieval of relevant information difficult and is one of the major deficiencies of current search engines. Problems concerning location-based data retrieval are well summed-up and solutions are proposed in the project of the Alexandria Digital Library [11].

2.4 Local Search Engines

The importance associated with local search capabilities by users and consecutively by the search engine industry is well illustrated by the history of the Google Local Search engine (<http://local.google.com/>, beta). It emerged from the 2002 Google Programming Contest winning project, the geocoding project of Dan Egnor (<http://dan.egnor.name/google.html>). The engine searches for address patterns on the Web sites and geocodes them using US Census TIGER datasets. While Google pioneered local search based on web site content analysis, other services provided limited search based on location before. Some providers profited of the availability of “Yellow pages” databases, geocoded to national street network datasets. In Australia, this sort of service is provided by Sensis (<http://www.sensis.com.au>).

During the drafting of the experiment, Yahoo launched its “local search” service for Australia, enabling to search by keywords and by location (usually entered as suburb name and street name combination). This proved to be very useful and provided high success rate of search.

Finally, a set of open source tools of the project `MobileMaps.com` enables anybody disposing of appropriate geographic data to create his own location based search engine, analysing the Web and mapping it to the geographic datasets.

As mentioned in [9], search engines such as Google or Yahoo cannot use the information contained on the Web site unless it has a well defined pattern or contains specific, clearly advertised location based keywords. Contact and location information hidden in databases and stored as dynamic content cannot be used by those engines.

2.5 Georeferencing the Web Content

To enhance the quality of location-based search on the Web, explicit georeferencing of the content will be necessary. The ultimate goal is represented by the Geospatial Semantic Web, and technologies enabling it are slowly emerging. Ontologies specified using the new WWW Consortium’s standard for ontology encoding, Web Ontology Language (OWL, <http://www.w3.org/2004/OWL/>) and encoded in the Resource Description Framework (RDF, <http://www.w3.org/RDF/>, <http://www.w3.org/2003/01/geo/>), will allow precise and machine readable coding of the resources located on the webpages and will provide a reliable mean to assess its content, including its geography [8]. This is the central focus of collaborative mapping communities, willing to provide citizens with additional

information about their neighborhood, public policies and local decision making processes (<http://space.frot.org/>).

The Semantic Web technologies are still far from wide adoption. Therefore, concurrent initiatives are supported to provide bridging capabilities to the current Web. People need to annotate and search for localised photographs [12, 13], blog entries, company websites [14]. Geographic HTML tags, enable some search engines to identify the location described by the Web resource. However, there is no major standard technology implemented by search engines, making wider adoption of geographic tagging impossible. The ICBM meta tags advocated and supported by the GeoURL project (www.geourl.org, deprecated) seems to be still very popular and supported by other services and softwares (<http://www.mapbureau.com/>, <http://www.a2b.cc/>).

Different focus of GIS professionals and general public creates a gap and inhibits coordinated development. The Open Geospatial Consortium's Geography Markup Language (GML) targeted on GIS professionals is too complex to be adopted by general public. Further concerns arise from the developers community, lacking the support for richer semantics and OWL/RDF annotation of GML. An unofficial translation is, however, available (<http://loki.cae.drexel.edu/~wbs/ontology/ogc-gml.htm>).

Finally, formal semantics for communicating results of geospatial queries need to be developed. First attempts to enable automated communication of routing instructions are proposed in [15], but semantically enriched encoding is still missing.

3 Wayfinding With Landmarks

Navigation is a complex process, requiring the mover to perceive the environment and relate his observations to the instructions provided prior or during the navigation. While navigating, the mover compares his perception of the world to his mental model of the environment. Routing instructions usually consists of a descriptions of specific reorientation points connected by route segments, along which the mover navigates. Michon and Denis [16] suggest that landmarks are used as assistance to movers to construct a mental representation of unfamiliar environments prior to their visit. They also show that the general need for landmarks increase at specific points on the route, specifically at the decision points where reorientation takes place. Furthermore, navigators find it useful to include landmarks also along longer segments, where those act as routemarks and are used to confirm the correctness of the path used. The frequency of routemarks is highly dependent on the mode of transport and personal preferences of the mover.

Another reason to include landmarks in route directions is the human capacity to refer to unknown locations by referring to close, familiar reference points [17], or to points that are salient with comparison to their neighborhood. Objective aspects of saliency—saliency not dependent on previous individual experience or on specific aspects of human context, but only on the visual prominence

of their façades—have been explored by Nothegger, Raubal and Winter [18], and applied in [19]. The results are in compliance with human selection of landmarks.

Current navigation Web-service do rarely provide landmarks as part of the instructions. We can assume that one of the reasons lies in the bad accessibility of complete data and the resulting costs of databases necessary for such a service. Instead, the directions provided to pedestrians have exactly the same form as those provided to movers using different modes of transport—i.e. using street names and distance indication. The only adaptation is at the level of street networks, where driving restrictions do not apply, and the movers speed is adjusted. Routing directions based uniquely on metrics (be it time or distance) are however unnatural to pedestrians, as their perception of distance can be highly distorted by external factors (difficulties in estimating distances due to lack of odometer, crowds on the street, traffic lights causing delays, . . .).

Personal navigation related tasks are one of the most common spatially related problems solved by the general population in everyday life. It is therefore an important use-case that should be supported by the geospatial Semantic Web. However, if it is to be supported by landmarks, the spatial coverage and the level of detail of the Web resources needs to be high.

4 Experiment Description

4.1 Web-Assisted Retrieval of Landmarks

To assess the spatial coverage of real world phenomena by the Web, we devised a simple case study, that could be automated as a Web-service. The test simulates the actions which could be implemented by a Web service. A route was generated using an on-line routing service. The areas, streets and locations along the path were manually searched using Web-search services, with location-based search options. An additional search was performed on places where we expected relevant information (city council databases, road repair databases, directories). Finally, search results were filtered, and the final set was used to generate route instructions including landmarks and routemarks. The resulting set of instructions was tested in human subject testing in a navigation exercise. Fig. 1 presents a flowchart of our case study.

4.2 Testing Environment

Our experiment was performed in Melbourne, Victoria, the second largest city in Australia. Melbourne has a high population of Internet users and density of Internet hosts. We can assume that this fact will lead to a good description of Melbourne by Web resources, and that the coverage of the City Business District (CBD) will be higher than at the outer parts of the centre. Therefore, a route covering two parts of the city, with differing urban settings, was planned for comparison. The navigator, a tourist, was trying to find his way from the Royal Melbourne Hospital through to the Windsor hotel (where she wanted to take lunch), and then to the Central station on Flinders street.

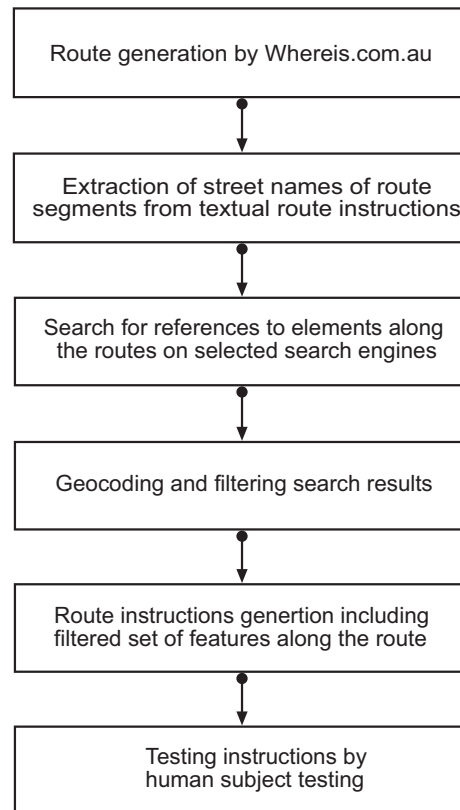


Fig. 1. Experiment flow chart

4.3 Search Method and Tools

We have used a specialised Australian mapping Web service www.Whereis.com.au to generate shortest path routing instructions for a pedestrian navigating in the CBD. Names of the streets were extracted from the routing directions provided by the service.

Consecutively, the street names were searched on the Web using search engines Google (<http://google.com.au>, with “Australian sites only” option selected), Yahoo (and its Local Search <http://localsearch.yahoo.com.au>), Sensis (<http://www.sensis.com.au>) and Mooter (www.mooter.com), to get all the information referenced to the search keywords. Where possible, images of the location described where searched, mainly using Google Image search. Additional searches were performed in specialised databases and directories (Victorian Heritage Register—<http://www.doi.vic.gov.au/doi/hvolr.nsf>, Citysearch—<http://www.citysearch.com.au>). These sources present the “hidden Web” content — databases and dynamic Web pages, not searchable by current search engines.

Currently, only humans can access and query those resources. Wider adoption of Web service discovery and capabilities declaration for automated processing on the Web is still lacking.

If the search result did not provide a map positioning the resource, its address was geocoded on WhereIs, to get its exact location along the path. The localisation of the location information on the site was noted for later analysis of Web designer ontologies.

4.4 Processing Results for Route Instructions

When selecting elements for inclusion in the instructions, general rules about landmarks applied [16]. Using Web content for navigation inherently increases the uncertainty of quality of service. Therefore, a higher density of elements was included in the instructions, considering the possibility that some of the resources may not be present or salient enough. The resources were filtered according to the following rules, in the respective order:

1. Every decision point should be covered by a landmark, if a resource is available;
2. On segments longer than one block, every block should contain a routemark if available. Support at decision points where no change of direction is taken may be important to the mover.

For locations with redundant or excessive amount of hits, we have filtered the set of features, resulting from general theories of landmark inclusion in navigation instructions, and from the specific characteristics of the Web. These is an extensive set of rules, and only a subset was applied in our case study, as the subset produced already complied to the requirements stated above:

1. Only one landmark at each decision point should be included. The selection followed additional rules.
2. Prominence of the element in the Web—we assumed that a higher ranking in browser engines and a higher amount of Web resources (including photographs) pointing to the same element describes a more important feature of the real world. We expected that the visual prominence of this feature will be proportional.
3. Position with relation to the crossing. Elements found on the opposite site of the crossing were given preference, if possible. Previous experiences suggest that features located on the same side of the road are often difficult to notice. Human need some distance to be able to perceive façades. The same applies to crossings, where the façade included in the instructions can easily be located around the corner.
4. Resources pointing to specific flats were excluded (dentist offices,...), as the visibility of those elements is generally low.
5. Time since last update of the Web-resource. The older the resource, the less reliable its content is. More recent resources were given priority.

most prominent image found on the homepage of the resource was used, based on manual selection. Image processing algorithms can be devised to analyse the set of images found, and select those containing features appearing with the highest frequency. Names of landmarks with associated images were underlined in the instructions, to indicate a linked information.

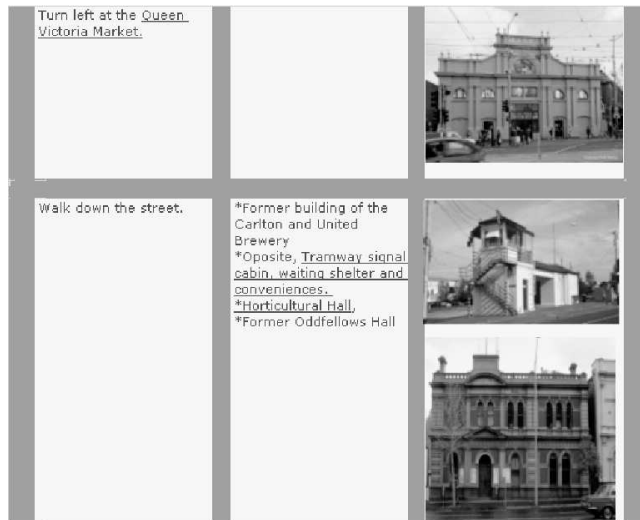


Fig. 3. Routing instructions with landmark references

Instructions in the set of directions directions used a unified terminology, but a consistent taxonomy of referencing landmarks along routes and in routing instructions is needed for automated navigation systems.

4.5 Testing Route Instructions Usability

The usability of the generated directions was assessed by a small human subject testing experiment. It was limited in extent and serves as proof-of-concept and provides feedback for the preparation of a larger experiment. Testers were guided to the start location, and from there navigated to the end of the route according to the instructions. Testers did not know the target of their route, and were not allowed to read further in the instructions.

This experiment tested the usability of the instructions and the quality of landmarks included. Features found on the Web were tested for presence (evaluation of obsolete or uncertain content of the Web), correctness of description (image of façade from correct angle, name of the resource matching the sign board on the façade), saliency and availability on important points of the route.

5 Outcomes

5.1 Webmasters' Ontology of Location Information

References to locations described by Web resources are inserted in the pages in a wide variety of ways. This reflects the differences in Web-masters ontologies, and a focus on a small range of use of location information. Among the Web resources found, most were referenced to locations by postal address patterns, in order to provide mail contact information or to communicate the location of offices, shops, etc. The way the information was embedded in the page itself was highly inconsistent. Keywords such as "Contact us", "Address", "Location", "Getting There" were most frequent. Bare location information in page headers or under a heading referring to a wider area (Melbourne CBD, Melbourne Metro), in case of store chains, was not unusual either. This may not be a problem for human interpretation, but pose problems to automated parsers, mainly in case of multiple addresses on a single page.

Even if the data about a specific landmark are available on the Web, their discovery can be complicated or impossible. Large compounds of buildings (i.e. university campuses) are often referenced by a single mail address, which makes the discovery of relevant information on the Web difficult. Web designers reference features as punctual objects, and neglect the coverage (extent and layout) of the features described, as well as its relation to surrounding features (road network, adjacent buildings). Such simplification can make the discovery of relevant objects impossible.

References to maps were common on sites of local search engines and directory services, with high unification of referencing keyword "Map". A Melbourne specific keyword referring to "Melway" map, the most common street directory of Melbourne was frequently found, indicating the map quad where the resource is located.

The richest location information about a specific resource was available on the pages of the Victorian Heritage Register, including "Spatial Information"—the latitude/longitude coordinates of the resource (again, a one dimensional reference), as well as its description ("Item Categories") and photographs. These data are stored in databases and are not accessible by standard Web search engines, unable to search the hidden Web.

5.2 Web Content Quality

This study consider the quality of the Web content for navigation purposes as its fitness-for-use. It includes aspects of availability (at specific points of the route, as indicated in the filtering criteria) and reliability (presence in reality with attributes extracted from the Web). Reliability of the Web resources covers problems as correct resource name, georeferencing (accuracy and coverage area definition), relevance of photographs.

Web Resource Availability:

- The availability of Web resources covering our testing area was good enough along most of the route and provided features that could be used as landmarks, except of two decision points that were not covered. In one case, this could be solved with a reference to the last routemark before the decision point, located in the same block. The second case could not be solved satisfactorily. A complex roundabout, connecting a complicated network of roads, was not described by any Web resource. Providing traditional routing instructions through references to street names, or by providing a numeric reference to the exit road seems to be the only possibility for such complex structures. No Web resources allowed us to include landmarks at this specific point.
- Not all blocks of the route segments were saturated with routemarks. This did not cause any problem to the navigators and the longer segments contained a satisfactory amount of routemarks for confirmation of turns at previous decision points.
- The fact that search engines list all resources for free, Web-site designers often try to get high rankings by publishing excessive and misleading keywords. This leads to a high proportion of irrelevant “spam” in the search results.
- The nondeterministic nature of the assessment of Web resources availability is further complicated by the nature of the search - the complexity of Web data storage and the large number of resources to search. It is impossible to assess the total amount of data searched and determine the proportion of data omitted.

Web Resource Reliability:

- Only one of the landmarks included in the instructions pointed to a non-existent feature, a freshly demolished building. The Web resource indicated that it describes the “former” building of the Carlton and United Brewery.
- The interpretation of “former” is ambiguous. It may also design a change of function, attribute, as it was the case of the “Former Oddfellows Hall”. A difference in the building’s name (ODDI Hall), led to a confusion of the mover. The new name was however indicated on the original Web page, but not included in the instructions, as far as it was located deep in the textual description of the building. The strict rules attempting to simulate an automated process did not allow its inclusion.
- Photographs of buildings are the source of a number of problems: views of façades taken from a different viewpoint, hidden by vegetation, or with obsolete content (façade after restoration, new sign board) and thus providing an unrecognisable image of a landmark are common. The Semantic Web and its capability to annotate photograph with rich attributes describing the content, location of the subject, date, time and direction in which the photograph was taken will help to partially overcome it.

- Incorrect insertion of photographs in Web-pages are a special problem—flipped photographs present a source of confusion for movers and automated image content extraction techniques.
- The content of Web resources used for the experiment was surprisingly up-to-date. This demonstrates one of the crucial assets of the Web—the maintenance of the contents is done by its originator, which also shares the interest to provide correct information to the user. The reliability of the content of Web directories is even higher, as the subjects are paying for being listed, and seek to assure that their contact data are correct. Web pages which pointed to nonexistent locations, such as sites of closed/renamed restaurants and shops, were not present found in our case, which may be just a coincidence. In general, a percentage of search results, points to nonexistent resources, often cached by search engines.

5.3 User Satisfaction

The purpose of the experiment was to assess the Web content for purposes of navigation with landmarks. Therefore, in our test purely landmark based instructions were used. It is clear that in real application not only landmarks (environmental features) but also other specific components (delimiters) should be used, as stated by Allen [21]. The deliberate choice to exclude references to metrics and street names influenced user satisfaction with the route directions. Still, this was necessary in order to assess the fitness-for-use of the references retrieved. Therefore, our basic criterion for user satisfaction lied in the testers' tendency to lose the correct route. One point of the route were identified as problematic by all participants. The lack of resources describing the layout of a complex junction—roundabout was not possible to solve. This signifies that simple substitution of metrics by landmarks in traditional route directions does not automatically lead to satisfactory results. Directions need to be modified to provide optimised information to the user and his context. Landmarks enhance route descriptions with more user friendly and natural references to the environment and add value to traditional instructions, but in complex environments both approaches need to be combined for full user satisfaction.

Further complaints focused on the issues mentioned above in Sect. 5.2, related to the reliability of landmark descriptions extracted from the Web—feature name, quality of the associated photograph, and mainly the saliency of the landmark in the environment. The assessment of the saliency of the landmark prior to its inclusion in the directions is crucial to the resulting quality of the whole set of instructions.

5.4 Problems and Open Questions

Features along the route, mostly routemarks, can be difficult or impossible to find, as they are referenced to auxiliary or parallel streets of the path, in case of larger buildings. As in our experiment we do not have direct access to the

street network database, the retrieval of those streets would require image interpretation at the level of the map provided by the Web service. This situation is similar to that of a third-party service provider, not having direct access to additional data from the source. Still, at simple street junctions where decisions are taken, information about all corner buildings may be retrieved. This enables to search for landmarks at decision points. However, complicated junctions and roundabouts are still problematic to handle.

To increase the cognitive plausibility of the resulting set of instructions, it is necessary to adapt the selection of landmarks to fit the context of the navigator. Thus, the area of user context definition, determination of its basic elements influencing human perception of the environment, and the consequent formalization of those elements based on ontology of context requires additional research effort. Consecutively, adapting the granularity of instructions at different parts of the route is crucial to provide more user friendly instructions, fitting the user's perception and knowledge of the environment

Apart of technical problems with georeferencing Web resources (Sect. 2.5), yet another complication is caused by the different "purpose" of most of the georeferencing information contained in the current Web. Web creators rarely consider navigation with landmarks as a use-case when designing their products. Their ontology is radically different from the ontology of a route planner. Therefore, one of the basic requirements for the (Geospatial) Semantic Web is the mapping between different ontologies and support for spatial operations, including natural language structures in the queries (near, close to, after).

Finally, the formalisation of satisfying route instructions is at least as problematic as the retrieval of route segments and environmental features. Additional research is needed to standardise instruction taxonomies where landmarks need to be referenced in a specific way. Similar taxonomy was already developed for route directions symbolisation [22].

6 Conclusions

For lack of previous estimates of Web content richness prior to the study, an objective to find at least a half of all the geocoded resources in the reality was set. A set of features evenly distributed along the path, including endpoints and a majority of decision points covered, was the goal. The urban environment of an industrial and developed city showed to be well described by the Web content and enabled to rich these goals. The available content of the Web was usable for the purpose of the wayfinding application.

The experiment showed that the Web is exploitable as a source of spatial data for services requiring relatively high level of input detail—landmark-based urban navigation services, for example, even if the quality of its content is still uncertain. Purely landmark based Web navigation services can still not be developed. The Web can however provide elements that can well complement and enhance current routing directions based on street names and metrics.

The case-study provided insight in the capabilities and the coverage of the Web. The advent of the Semantic Web will mean a major enhancement in its usability. Geographic analysis will be possible and geocoding queries based on fuzzy queries (neighbourhood, proximity) will enhance future location and navigation services. This will necessarily be linked to advanced geocoding and data-mining algorithms. My future work will focus on specific aspects important for improvement of those services: the selection and formalisation of criteria applied to filtering the resources found, focused on selecting features with relevance to the navigator's context.

Acknowledgments

The author would like to acknowledge the CRC for Spatial Information for supporting the project. I would also like to thank Stefanie Andrae for her valuable help with the case study. Acknowledgments belong also to the participants of the human subject testing.

References

1. NUA: How Many Online? Available on: http://www.nua.com/surveys/how_many_online/index.html (2002) Visited on August, 21st 2004.
2. Zook, M.: Connected is a Matter of Geography. *netWorker* **5** (2001) 13–17
3. Zook, M.: Internet Users Worldwide. Available on: www.zooknic.com (2001) Visited on August, 21st 2004.
4. Shiode, N., Dodge, M.: Using GIS to Analyse the Spatial Pattern of the Internet in the United Kingdom. In: GIS Research UK 6th Nation Conference, Edinburgh, UK (1998)
5. Dodge, M., Kitchin, R.: Charting Movement: Mapping Internet Infrastructure. In Hanley, R.E., ed.: *Moving People, Goods, and Information in the 21st Century*. Routledge, New York (2004) 159–185
6. Berners-Lee, T., Hendler, J., Lassila, O.: The Semantic Web. *Scientific American* **184** (2001) 34–43
7. Egenhofer, M.J.: Towards the Semantic Geospatial Web. In Voisard, A., Chen, S.C., eds.: *ACM-GIS 2002*. Volume 6., McLean, VA (2002)
8. Hiramatsu, K., Reitsma, F.: GeoReferencing the Semantic Web: Ontology Based Markup of Geographically Referenced Information. In: Joint EuroSDR/EuroGeographics workshop on Ontologies and Schema Translation Services, Paris, France (2004)
9. Young, R.: Google Local Search and the Impact on Natural Optimization. *WebPronews.com* (2004)
10. Buyukkokten, O., Cho, J., Garcia-Molina, H., Gravano, L., Shivakumar, N.: Exploiting Geographical Location Information of Web Pages. In: *Proceedings of Workshop on Web Databases (WebDB'99)*. (1999)
11. Goodchild, M.: The Alexandria Digital Library Project. *D-Lib Magazine* **10** (2004)
12. Spinellis, D.: Position-annotated Photographs: A Geotemporal Web. *IEEE Pervasive Computing* **2** (2003) 72–79

13. Toyama, K., Logan, R., Roseway, A., Anandan, P.: Geographic Location Tags on Digital Images. In: MM03, Berkeley, California, USA, ACM (2003)
14. elasticspace: Spatial Anotation Projects. Available on: <http://www.elasticspace.com/index.php?id=35> (2004) Visited on October, 10th 2004.
15. Open GIS Consortium, I.: OpenGIS Location Services (OpenLS): Part 6 - Navigation Service. Technical Report OGC 03-007r1, Open GIS Consortium (2004)
16. Michon, P.E., Denis, M.: When and Why Are Visual Landmarks Used in Giving Directions? In: Spatial Information Theory. Foundations of Geographic Science. International Conference COSIT 2001. Lecture Notes in Computer Science, Morro Bay, CA, USA, September 2001, Springer-Verlag Berlin; Heidelberg; New York (2001) 292–305
17. Tversky, B.: Cognitive Maps, Cognitive Collages, and Spatial Mental Models. In Frank, A., Campari, I., eds.: Spatial Information Theory: A Theoretical Basis for GIS, COSIT '93. Volume 716 of Lecture Notes in Computer Science., Springer, Berlin (1993) 14–24
18. Nothegger, C., Winter, S., Raubal, M.: Computation of the Saliency of Features. *Spatial Cognition and Computation* 4 (2004) 113–136
19. Elias, B., Brenner, C.: Automatic Generation and Application of Landmarks in Navigation Data Sets. In: Spatial Data Handling, SDH 2004, Leicester (2004)
20. Raubal, M., Winter, S.: Enriching Wayfinding Instructions With Local Landmarks. In Egenhofer, M.J., Mark, D.M., eds.: Geographic Information Science. Volume 2478 of Lecture Notes in Computer Science. Springer, Berlin (2002) 243–259
21. Allen, G.L.: Principles and Practices for Communicating Route Knowledge. *Applied Cognitive Psychology* (2000) 333–359
22. Klippel, A.: Wayfinding Choremes. In Kuhn, W., Worboys, M., Timpf, S., eds.: Spatial Information Theory: Foundations of Geographic Information Science. Conference on Spatial Information Theory (COSIT). Lecture Notes in Computer Science, Springer, Berlin (2003) 320–334