

Acquisition and Relevance of Geotagged Information in Tourism

Astrid Dickinger^a, Arno Scharl^a, Hermann Stern^b
Albert Weichselbraun^c and Karl Wöber^d

^a Department of New Media Technology
MODUL University Vienna, Austria
astrid.dickinger@modul.ac.at, arno.scharl@modul.ac.at^b

Knowledge Management Institute,
Graz University of Technology, Austria
hermann.stern@tugraz.at

^c Institute for Information Business,
Vienna University of Economics and Business Administration, Austria
albert.weichselbraun@wu-wien.ac.at

^d Department of Tourism and Hospitality Management
MODUL University Vienna, Austria
karl.woeber@modul.ac.at

Abstract

In the case of tourism applications, it is particularly evident that geography is emerging as a fundamental principle for structuring Web resources. Recent improvements in semantic and geographic Web technology, often referred to as the *Geospatial Web*, acknowledge the relevance of adding location metadata to existing databases and accessing the vast amounts of information stored in these databases via geospatial services. This paper outlines the acquisition of geospatial context information, describes usage scenarios and real-world applications in the tourism industry, and presents an automated software tool for annotating large collections of Web documents automatically. The quality of this tool is tested based upon Web pages from the Austrian National Tourism Organization. Initial results are encouraging and help define a roadmap for further improving the automated tagging of tourism resources.

Keywords: geotagging, text mining, information retrieval, tourism information systems.

1 Introduction

Travel decisions require a large amount of information and encompass a number of information search, evaluation and integration tasks and activities (Hwang et al, 2006). Due to the nature of the service industry with intangibility of products and in particular in travel, consumption at a later point in time, information search has been a critical factor in the decision process in order to reduce travellers' risk (Lewis & Chambers 2000). The Internet emerged as a medium that provides travellers with relevant information and is slowly emerging as a booking platform (Gursoy & McLeary, 2004; Klein, Köhne & Öörni, 2004).

For travellers location and the destination are of particular relevance. Typically questions in tourism are where is the airport, where is the castle, where are nice restaurants. This indicates that providing information on a map that allows the searcher to immediately filter the relevant spots can drastically improve the information search experience. While the Internet provides some location indicated information, it does not, however, do so to the extent desirable for tourists. However, there is the emerging trend of including maps in web portals. Furthermore, the geospatial web, representing the integration of geographic and semantic information will fuel this development (Scharl & Tochtermann 2007). A simple Google search for *hotel* and *Munich* for instance yields just below three million results. The tourist searching for information now has to identify the pages that in fact deal with hotels and above that are located in Munich. Carrying out the same search on Google Maps arrives with 26,568 results. These are located on the map of Munich and easily found marked with a flag. This illustrates the usability of the system and the improvement of the actual search results.

The aim of this study is to present and discuss methods for automatically creating geotagged knowledge repositories (Scharl, 2007) in the tourism domain, to evaluate the accuracy of the resulting annotations, and to provide guidelines for further improvements of geotagging components. The article proceeds as follows. First manual and automated forms of tagging are presented, and then tourism applications of geotagged information are presented to provide some real world examples. The next section presents the automated tagging procedure, followed by an evaluation. The paper closes with conclusions and future research avenues.

2 Geo-Based Information in the Tourism Industry

There are attempts in the travel and tourism industry to use location information to provide better services to travellers. This type of information is primarily exploited for location-based services using Global Positioning Systems (GPS). One such example is a tourist guide for travellers along the river Danube (Dickinger & Zins 2006). Based on the location of the ship the travellers are provided with video clips, information on events, the surrounding area, pictures, and historical information on an interactive map. Location-based services, rather innovative even in the maritime market (Maglogiannis, Kormentzas and Panagiotarakis, 2004a), use the current position of the ship to provide richer information to the customer. These services were primarily push based, the user is located and the content is pushed to a device such as computer, mobile phone or personal digital assistant, to name a few. Other services allow the user to search information in the context of the location he or she defines. This means that maps are provided and the information is location indicated or geotagged.

A geo-based information platform for tourism includes *TrustedPlaces.com* which provides recommendations by users. They can inform their peers about restaurants, pubs, clubs points of interest, sights etc. This information is then linked with Google maps to locate the specific place. *Tupelo.com* has a similar concept with users posting their experiences with restaurants, cafes and bars. However, the initial aim was not to provide information for tourists but people who live nearby.

The project *Personal Experience with Active Cultural Heritage (PEACH)* enhances cultural heritage appreciation by using moves and gestures of visitors to a museum as implicit input for guiding him/her through the building and for hinting the person to specific paintings (Stock et al. 2004). The system combines geo-based information (the location of the visitor as well as the paintings) with persuasive communication tools in order to provide personalized and anticipated content for museum visitors.

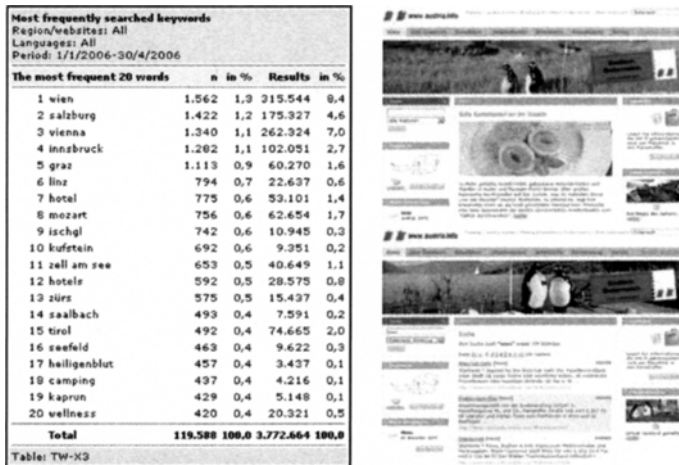


Fig. 1. Top 20 search terms entered on www.austria.info between January and April 2006 (left); screenshots of the domain-specific search engine (right).

The *Austrian National Tourist Office* (www.austria-info.at) is continuously improving technologies for human-to-machine communication. A domain-specific search function was introduced in 2006, for example. A comprehensive analysis of user inquiries for information on all Austrian provincial Web sites simultaneously hints at the potential of geo-enriched information and search features, as the vast majority of search terms entered by the users refer to notations of Austrian regions (see Figure 1). A different platform in collaboration with the *Austrian National Tourist Board* is Coolaustria.com. Its aim is to provide a portal for users that want to post their holiday experience, photos, and recommendations of places in Austria. Above these tips and information, also the travellers are shown on the map. It is possible to contact them in case they are experts for a trip or activity one plans to do.

3 Identifying and Annotating Geospatial Context Information

Efforts are under way to annotate existing tourism resources with geospatial metadata. This process, usually referred to as *geotagging*, assigns geospatial context information, ranging from specific point locations to arbitrarily shaped regions. Geotags are crucial to build advanced tourism information systems and have a wide range of possible applications. They can be used to identify and match users with similar interests, facilitate their communication and collaboration, and provide them with intuitive

search support as shown in Figure 2 – not only considering topical similarity, but also on the location of an information object relative to the user’s actual location. Examples include *Jobloft.com*, which enables users to search for retail jobs, food services jobs and hospitality jobs in their vicinity, and *Housingmaps.com* to find out the location of apartments and houses via Google Maps.



Fig. 2. List of articles related to the term ‘vacation’ within the geo-referenced archive of the IDIOM Media Watch on Climate Change (www.ecoresearch.net/climate)

Sources of Geospatial Information. Different sources of geospatial context information for annotating Web resources often complement each other in real-world applications (McCurley, 2001):

- Annotation by the author, manually (Daviel & Kaegi, 2003) or through location-aware devices such as car navigation systems, RFID-tagged products and GPS-enabled cellular handsets. These devices geotag information automatically when it is being created. Tagging by the author allows the inclusion of different views and associations with specific text fragments or photos. Within small virtual communities, collaborative manual tagging is still predominant.
- Determining the location of the server – by querying the *Whois* (www.whois.net) database for a hotel’s domain registration, monitoring Internet traffic, analyzing the domain of a Web site for additional cues (e.g. www.austria.info), or by using dedicated services such as the *Geo IP Tool* (www.geoiptool.com), which returns longitude and latitude of specific IP addresses.
- Automated annotation of existing documents. The processes of recognizing geographic context and assigning spatial coordinates are commonly referred to as *geoparsing* and *geocoding*, respectively.

Once geospatial context information becomes widely available, any point in space will be linked to tourism-related information such as hotel descriptions, movie programs, exhibitions, or personal stories and preferences. Even locative spam will become a common phenomenon (Erle, Gibson, & Walsh, 2005) with the widespread introduction of location-based services.

At present, however, many initiatives to provide geographically referenced tourism information suffer from the chicken and egg problem, wishing that existing content was retrofitted with metadata (McCurley, 2001). This “capture bottleneck” results

from the beneficiaries' lack of motivation to devote the necessary resources for providing a critical mass of metadata (Motta, Shum, & Domingue, 2000). Geotagging projects are no exception. Acknowledging calls to automate the semantic annotation of documents (Benjamins, Contreras, Corcho, & Gómez-Pérez, 2004; Domingue & Motta, 2000), the following sections focus on the third category, the automated geoparsing and geocoding of existing tourism resources – information published by national tourism offices, regional information, hotel descriptions, blogs of travellers during and after their holidays, and other types of unstructured textual data found on the Web.

Geoparsing. All human artifacts have a location history, which commonly includes a creation location and a current location (Spohrer, 1999). Given the availability of geographic references, Web applications can map the whole life cycle of such artifacts. Tourism resources are particularly rich in explicit or implicit geographic references. This includes references to physical features of the Earth's surface such as forests, lakes, rivers and mountains, and references to objects of the human-made environment such as cities, countries, roads and buildings (Jones, Alani, & Tudhope, 2001). Addresses, postal codes, telephone numbers and descriptions of landmarks also allow us to pinpoint exact locations (Ding, Gravano, & Shivakumar, 2000; McCurley, 2001). At least 20 percent of Web pages contain easily recognizable and unambiguous geographic identifiers (Delboni, Borges, & Laender, 2005). Web-based tourism resources are particularly rich in such identifiers – e.g. discussing a destination, reviewing accommodation options, or discussing a region's attractions and highlights.

Identifying and ranking spatial references by semantically analyzing textual data is a subset of the more general problem of *named entity recognition*, which locates and interprets phrasal units such as the names of people, organizations and places (Cowie & Lehnert, 1996; Weiss, Indurkha, Zhang, & Damerou, 2005). As with most named entity recognition tasks, false positives are inevitable – e.g., documents that quote addresses unrelated to the their actual content (Morimoto, Aono, Houle, & McCurley, 2003). Ambiguity, synonymy and changes in terminology over time further complicate the geoparsing of Web documents (Amitay, Har'El, Sivan, & Soffer, 2004; Kienreich, Granitzer, & Lux, 2006; Larson, 1996). Identical lexical forms refer to distinct places with the same name (VIENNA refers to the capital of Austria as well as a town in Northern Virginia, USA) or have geographic and non-geographic meanings: TURKEY (large gallinaceous bird; bi-continental country between Asia and Europe), MOBILE (capable of moving; city in Alabama, USA), or READING (processing written linguistic messages; town in Massachusetts, USA). Geoparsers also need to correctly process references to identical or similar places that may be known under different names, or may belong to different levels of administrative or topographical hierarchies (Jones et al., 2001).

Geocoding. Once a location has been identified, precise spatial coordinates – latitude, longitude and altitude – can be assigned to the documents by querying structured geographic indices (gazetteers) for matching entries (Hill, Frew, & Zheng, 1999; Tochtermann et al, 1997). This process of associating documents with formal models is also referred to as document enrichment (Domingue & Motta, 2000; Motta et al.,

2000). Examples of formal geographic models are the Geographic Names Information System (geonames.usgs.gov), the World Gazetteer (www.world-gazetteer.com), the classifications of the United Nations Group of Experts on Geographical Names (unstats.un.org/unsd/geoinfo), the Getty Thesaurus of Geographic Names (www.getty.edu) and ISO 3166-1 Country Codes (www.iso.org/iso/en/prods-services/iso3166ma/index.html).

While simple gazetteer lookup has the advantage of being language-independent, advanced algorithms consider lexical and structural linguistic clues as well as contextual knowledge contained in the documents; e.g., dealing with ambiguity by removing stop-words, identifying references to people and organizations (Clough, 2005) and applying contextual rules like “single sense per document” and “co-occurring place names indicate nearby locations”. Each identified reference is assigned a probability $P(\textit{name}, \textit{place})$ that it refers to a particular place (Amitay et al., 2004). The location that receives the highest probability is then assigned a canonical taxonomy node such as EUROPE/AUSTRIA/VIENNA; 48°14' N, 16°20' E.

4 Implementation of a Geotagger

The geotagger module outlined in the following section annotates documents, sentences, or even single terms with tags describing their geographic context. In contrast to methods which determine only the source geography of a Website – i.e., the location of the content provider – based on the host name or its IP address, a geotagger provides additional information about geospatial references in an article's content. For example, a promotional Web document describing the beautiful lakeside of the Wörthersee will be tagged with a geographic focus on Carinthia, Austria.

Leveraging the latest version of the GATE framework – the General Architecture for Text Engineering; <http://www.gate.ac.uk> (Cunningham, Maynard, Bontcheva & Tablan 2002) – the part-of-speech tagger identifies locations by using a gazetteer – a dictionary containing all relevant terms, referring to a particular location. The tagger's performance and accuracy is highly dependant on the number of gazetteer entries. Thus, the best strategy is to provide localized versions of the gazetteer covering different regions with customized scales and accuracy, and the functionality to switch between them. For a tourism use case in Austria the result may include Austrian villages, mountains, lakes, rivers and points of interest like castles, museums and other sights, however, in a world wide context only cities above a certain number of inhabitants are relevant. To keep the geotagger language independent, dictionaries in different languages are used. Common software applications select the gazetteer language and region by using a combination of ISO 639-1 (International Organization for Standardization 2002) language and ISO 3166 (International Organization for Standardization 2007) country codes. The gazetteer for Austrian locations in German corpora therefore is selected by using the code “de_AT”, “en_AT” in contrast refers to Austrian locations within an English speaking text.

Free geographical databases like GeoNames.org provide detailed information about geographic entities, including their names in different languages, their geographic

position (latitude and longitude), administrative divisions, populations, the location's timezone, etc. in an easily parseable format. Refining these data yields the language and region specific geo-database used in the geotagger.

Figure 1 describes the tagging process: The tagger identifies geo-locations based on the selected gazetteer. A focusing algorithm selects the most suitable locations for documents containing multiple possible geographic foci. Based on the information in the geo-database a tree is built, aligning locations hierarchically (for example Austria is a child of Europe, Vienna a child of Austria, etc.). This tree can have any number of levels, including for example states, areas, points of interest and so on.

Based on weights assigned to all locations present in the text, the weights of parent nodes are derived providing a list of geographic places sorted by their importance. Factors like the locations population, the number of occurrences in the text, and the tagger's default settings determine the assigned weights. Depending on the use cases, the tagger can return only the most relevant result and it's probability or a list of all possible matches (taking an article about a number of countries in the European Union as an example, the focus "Europe" is not satisfying – providing all found countries would do a quite better job).

The number of occurrences is only one indicator for the significance of the found location, in the case of ambiguities (two or more locations found the same number of times), the location tree is used to derive the most suitable candidate. At standard settings, one occurrence of Vienna and one of Graz deliver "Austria" as result, because that is the smallest location including both found terms. This behaviour, however, can easily be changed by modifying the algorithm's parameters. A more detailed description of the focusing algorithm can be found in Amitay, Har'El, Sivan, & Soffer (2004).




Fig. 3. Automated geo-tagging process

Publishing the tagger as Web Service running as a Tomcat Servlet (<http://tomcat.apache.org>) facilitates interoperability and easy integration into third

party applications. Most popular programming languages like Java, Perl, Python, PHP, C, C++, and even many content management systems as for instance Drupal offer this functionality. This setup also allows us to combine multiple geo-tagger installations into service clusters, providing higher throughput and availability.

The current implementation of the tagger is capable of annotating 1,000 text documents with a gazetteer comprising over 6,000 locations in 2:30 minutes on a Dual-Intel Xeon 3.2 GHz machine with two tagger threads. The tagger's performance can be improved considerably by clustering or using a higher number of tagger threads.

5 Evaluation of the Geo Tagger

Following the above procedure, a sample of tourism web sites from national and regional tourism organizations (NTOs, RTOs) were included. The webLyzard crawler (Scharl et al. 2004) was used to mirror and store the documents of the tourism sites and then automatically geo-tagged them. A random sample of 104 tagged documents were then chosen and analyzed. The data file included an URL of the target document, the automatically identified locations with latitudes, longitudes and the location name. Criteria to manually analyze the accuracy of the tagger were identified.

Each of the pages was visited and the actual city, village, or region the document is dealing with was identified. Furthermore, information provided on the web site such as subdomain, content of the title of the page, local area code, mentioning of phone numbers, frequency of the city name in relation to the amount of text that may hint at ways and indications to improve the tagging procedure was collected.

In a first step the general information of a city or region was collected, also the longitudes and latitudes of the real place were identified. The results indicate that 22% of the documents deal with cities, 78% with regions. This allowed for a comparison of the distance between the automatic tag and the actual location. The average distance between the real place and the automatic tag is 80km, for 50% the distance is only 44km. A total of 49% of the tags were correctly assigned to a city within a region, 34% of the documents were identified to deal with Austria and 17% showed the wrong place. A further analysis of the wrong classification showed that in the majority of cases (66%) the correct name of the place was not mentioned frequently enough. This is followed by 20% of the cases where two places were mentioned equally often and the wrong one was picked by the tagger. The last source of wrong classification is ambiguity of the location. In 15% of wrong classifications the name of the place has multiple meanings and could not be correctly identified.

As a next step the correct location was compared with the subdomain, this revealed that in 65% of the cases the subdomain matched the correct location. Furthermore, in 78% the title of the page matched the location and the imprint showed the correct location in 58% of the cases. However, it needs to be mentioned that the imprint might not be the ideal source for improvement of the tagger. The information mentioned there has to be used cautiously because often the web master or a company providing web services is mentioned, not the tourism office or the relevant sight/place. The area code and the place matched the actual location in 68% of the cases and the prefix was

correct in six out of ten cases. In addition, the pictorial elements on the web pages were included in the analyses. Eight out of ten times the picture matched the place.

The number of times the city was mentioned was measured (on average 4.3 times in one document) as well as whether the tag was correct. As expected a significant ($p < .05$) positive correlation of .25 between number of mentioning and correct tagging was identified. This implies that the more often a place is mentioned the more likely the tagger identifies the correct place. Analysis of the mentioning of the federal state and the accuracy of the tag yielded a correlation of .564 ($p < .01$). The same accounts for the region with a correlation of .384 ($p < .01$). A region is mentioned on average 3.4 times in one web document. A further analysis included the amount of text on a page in relation to the number of times a place is mentioned. The results show that there are on average 4,564 characters on a web page. A significant positive relation was identified between the frequency of the city name in relation to the text and the accuracy of the tag ($r = .239$; $p < .05$).

The results imply a couple of improvements for the tagger. Since the correct name of the location is included in the subdomain and the title of the page (86%) this information could improve the tagger. A total of 84% of correct tags would be reached. The inclusion of the area code and the prefix would improve this by 11% and arrive at a proportion of 95% of correct tags.

6 Discussion and Conclusion

The commonly foreseen intelligent mobile systems that are embedded, personalized and adaptive, and which anticipate user needs, require methodologies for filtering, weighting, and sorting of content. This paper has examined the opportunities and limitations of geo-tagging. The study tested an automated tagging procedure. The results show that the first implementation of the tagger works well but can be improved to arrive with even more accurate results. The inclusion of information from the website which is included in the crawling process can be used to arrive with an improved tagging procedure.

The complexity of today's information society is correlated with the overabundance of data caused by the Internet and other information technology accomplishments. In this sense, technology increases complexity as well as uncertainty for the users. However, it appears that technology is also the only means to reduce uncertainty, which implies even more information technology applications. This paper presented automatic geotagging as one of these approaches, which could help users to better manage this paradox development.

References

- Amitay, E., Har'El, N., Sivan, R., & Soffer, A. (2004). *Web-a-Where: Geotagging Web Content*. 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Sheffield, UK. ACM Press. 273-280.
- Benjamins, R., Contreras, J., Corcho, O., & Gómez-Pérez, A. (2004). Six Challenges for the Semantic Web. *AIS SIGSEMIS Bulletin*, 1(1), 24-25.
- Clough, P. (2005). *Extracting Metadata for Spatially-Aware Information Retrieval on the Internet*. 2nd International Workshop on Geographic Information Retrieval (GIR-2005), Bremen, Germany.
- Cowie, J., & Lehnert, W. (1996). Information Extraction. *Communications of the ACM*, 39(1), 80-91.
- Cunningham, H., Maynard, D., Bontcheva, K. & Tablan V (2002). GATE: A Framework and Graphical Development Environment for Robust NLP Tools and Applications. Proceedings of the 40th Anniversary Meeting of the Association for Computational Linguistics (ACL'02). USA, Philadelphia. <http://eprints.aktors.org/90/01/acl-main.pdf> {03-11-2007}
- Daviel, A., & Kaegi, F. A. (2003). *Geographic Registration of HTML Documents (IETF Internet-Draft, July 2003)*. Sterling: Internet Engineering Task Force. ecotroph.net/geopriv.
- Delboni, T. M., Borges, K. A. V., & Laender, A. H. F. (2005). *Geographic Web Search based on Positioning Expressions*. 2nd International Workshop on Geographic Information Retrieval (GIR-2005), Bremen, Germany. 61-64.
- Dickinger, A. & Zins, A. (2006). Adoption of Innovative River Cruise Information Systems. In M. Hitz, M. Sigala, J. Murphy (Eds.), *Information and Communication Technologies in Tourism 2006*, (pp 209-220). Wien: Springer.
- Dickinger, A. (2006). Mobile Datendienste im Tourismus. In R. Bachleitner, R. Egger, T. Herdin (Eds.), *Innovationen in der Tourismusforschung: Methoden und Anwendungen* (pp. 197-217). Wien: Lit.
- Ding, J., Gravano, L., & Shivakumar, N. (2000). *Computing Geographical Scopes of Web Resources*. 26th International Conference on Very Large Data Bases, Cairo: Morgan Kaufman. 545-556.
- Domingue, J., & Motta, E. (2000). PlanetOnto: From News Publishing to Integrated Knowledge Management Support. *IEEE Intelligent Systems*, 15(3), 26-32.
- Erle, S., Gibson, R., & Walsh, J. (2005). *Mapping Hacks - Tips & Tools for Electronic Cartography*. Sebastopol: O'Reilly.
- Gursoy, D. & McLeary, K.W. (2003). An Integrative Model of Tourists' Information Search Behavior. *Annals of Tourism Research* 31 (2): 353-373.
- Hill, L. L., Frew, J., & Zheng, Q. (1999). Geographic Names - The Implementation of a Gazetteer in a Georeferenced Digital Library. *D-Lib Magazine*, 5(1), <http://www.dlib.org/dlib/january99/hill/01hill.html> {03-11-2007}.
- Hwang, Y.-H., Gretzel, U., Xiang, Z. and D.R. Fesenmaier (2006) Information Search for Travel Decisions. In D.R. Fesenmaier, H. Werthner, K.W. Wöber (Eds.), *Destination Recommendation Systems* (pp 3-16). Wallingford: CABI.
- International Organization for Standardization (2002). Codes for the Representation of Names of Languages, International Organization for Standardization, http://www.iso.org/iso/language_codes. {03-11-2007}.
- International Organization for Standardization (2007). ISO 3166 Maintenance agency (ISO 3166/MA) – ISO's focal point for country codes, International Organization for Standardization, http://www.iso.org/iso/country_codes. {03-11-2007}.
- Jones, C. B., Alani, H., & Tudhope, D. (2001). Geographical Information Retrieval with Ontologies of Place. In D. R. Montello (Ed.), *International Conference on Spatial Information Theory: Foundations of Geographic Information Science (= Lecture Notes in Computer Science, vol. 2205)*, (pp 322-335). Berlin: Springer.

- Kienreich, W., Granitzer, M., & Lux, M. (2006). *Geospatial Anchoring of Encyclopedia Articles*. 10th International Conference on Information Visualisation (iV-06), London, UK. IEEE Computer Society. 415-419.
- Klein, S., Köhne, F., Öörni, A. (2004). Barriers to Online Booking of Scheduled Airline Tickets. *Journal of Travel & Tourism Marketing* 17 (2/3): 27-39.
- Larson, R. R. (1996). Geographic Information Retrieval and Spatial Browsing. In L. Smith & M. Gluck (Eds.), *GIS and Libraries: Patrons, Maps and Spatial Information* (pp. 81-124). Urbana-Champaign: University of Illinois.
- Lewis R. & Chambers R. (2000). *Marketing Leadership in Hospitality, Foundations and Practices*. New York: Wiley.
- Maglogiannis, I., Kormentzas, G. & Panagiotarakis, N. (2004a). Emerging Web-based Services for Ship Travellers. *Information Technology & Tourism* 7 (1): 23-31.
- McCurley, K. S. (2001). *Geospatial Mapping and Navigation of the Web*. 10th International World Wide Web Conference, Hong Kong, China. In V.Y. Shen and N. Saito (Eds.) Association for Computing Machinery. 221-229.
- Morimoto, Y., Aono, M., Houle, M. E., & McCurley, K. S. (2003). *Extracting Spatial Knowledge from the Web*. Symposium on Applications and the Internet (SAINT-2003), Orlando, USA. IEEE Computer Society Press. 326-333.
- Motta, E., Shum, S. B., & Domingue, J. (2000). Ontology-Driven Document Enrichment: Principles, Tools and Applications. *International Journal of Human-Computer Studies*, 52(6), 1071-1109.
- Scharl, A. (2007). Towards the Geospatial Web: Media Platforms for Managing Geotagged Knowledge Repositories. In A. Scharl & K. Tochtermann (Eds.), *The Geospatial Web - How Geobrowsers, Social Software and the Web 2.0 are Shaping the Network Society* (pp. 3-14). London: Springer.
- Scharl, A., Wöber, K.W. & Bauer, C. (2003). An Integrated Approach to Measure Web Site Effectiveness in the European Hotel Industry. *Information Technology & Tourism* 6(4): 257-271.
- Spohrer, J. (1999). Information in Places. *IBM Systems Journal*, 38(4), 602-628.
- Stock, O., Zancanaro, M. & Not E. (2004). Intelligent Interactive Information Presentation for Cultural Tourism. In O. Stock and M. Zancanaro (Eds.), *Intelligent Multimodal Information Presentation* (pp 95-111). Netherlands: Springer.
- Tochtermann, K., Riekert, W.-F., Wiest, G., Seggelke, J., & Mohaupt-Jahr, B. (1997). Using Semantic, Geographical, and Temporal Relationships to Enhance Search and Retrieval in Digital Catalogs. In C. Peters, C. Thanos (Eds.), *1st European Conference on Research and Advanced Technology for Digital Libraries (= Lecture Notes in Computer Science, Vol 1324)*, Berlin: Springer. 73-86.
- Weiss, S. M., Indurkha, N., Zhang, T., & Damerou, F. J. (2005). *Text Mining - Predictive Methods for Analyzing Unstructured Information*. New York: Springer.

Acknowledgement.

The IDIOM Project (Information Diffusion across Interactive Online Media; www.idiom.at) is funded by the Austrian Ministry of Transport, Innovation & Technology and the Austrian Research Promotion Agency within the strategic objective FIT-IT (www.fit-it.at). Section 3 is partially based on the introductory chapter of *The Geospatial Web – How Geobrowsers, Social Software and the Web 2.0 are Shaping the Network Society*. Eds. A. Scharl and K. Tochtermann. Springer: London (www.geospatialweb.com).