Map-Based Linking of Geographic User and Content Profiles for Hyperlocal Content Recommendation

Steven Verstockt[1(✉)], Viktor Slavkovikj[1], and Kevin Baker[2]

[1] ELIS Department - Multimedia Lab Gaston Crommenlaan, Ghent University–iMinds, 8 Bus 201, 9000 Ghent, Belgium
steven.verstockt@ugent.be

[2] Department of Geography, Ghent University–Rotuteyou, Krijgslaan 281, S8, 9000 Ghent, Belgium

Abstract. In this paper we describe a novel approach for map-based linking of users with content (and vice versa) based on their geographic profiles. The proposed technique facilitates hyperlocal content recommendation targeted to the user's geographic footprint. The generation of the geographic user profiles (GUP) is based on user-logged activity analysis. The result of this analysis is a heat map of the geographic keypoints where the outdoor activities were performed. For the geographic content profiling (GCP), we use the available geotags and perform address geocoding and geographic named entity recognition to extract additional locations from the media objects. In order to link the GCP to the GUP, and to be able to recommend the hyperlocal content that fits the user's current profile, heat map analysis is performed using geographic analyzing tools. The GEOprofiling demonstrator, which is evaluated on real activity profiles and different media types, shows the feasibility of the proposed approach.

Keywords: Recommendation · Geo-profiling · Geocoding · Heat map analysis

1 Introduction

Local media plays an important role in each of our lives. It is both functional, telling us what is going on in our direct environment, and emotional, helping us to feel like we belong to a local community. Furthermore, consuming media on the move has become a mainstream behavior for many of us, and more and more we expect media to be related to our current location. As such, it is not surprisingly that a huge increase is observed in the volume and usage of (online) local media. Traditional media, like newspapers, have observed this tendency, and start to focus more on localized content in their digital environments [1]. A similar evolution is seen on social media platforms, allowing people to connect to timely, relevant information at the hyper local level. Whoo.ly, for example, is a web service that provides neighborhood-specific information based on Twitter posts [2] and Field Trip is Google's hyperlocal recommendation engine for nearby sights and destinations.

S. Yamamoto (Ed.): HIMI 2015, Part II, LNCS 9173, pp. 53–63, 2015.
DOI: 10.1007/978-3-319-20618-9_6

Up till now, however, the way in which media platforms support (hyper)local content is still too much focused on large and static geographic areas/communities, i.e., too 'glocal', and does not take into account the specific and dynamic geographic footprint of the user. Similarly, the geographic description of the content is still too high-level and the annotation is still very often a manually labor-intensive process. Automatic geographic analysis at the micro level of the content, e.g. on scenes in a video or paragraphs in a text, is needed to create more valuable geographic content descriptions, which on their turn lead to more relevant queries.

The higher content scores on the two major dimensions of hyperlocality, i.e., geography and time, the more relevant the content becomes to the individual and the less it becomes to the masses. In order to be successful, hyperlocal content needs to be targeted to a specific user in a well-defined area and a specific time window. Nowadays, however, hyperlocal content recommendation is mostly based on static information from user profiles or device-based localization techniques. Furthermore, geographic annotation on the content level is mostly restricted too some geographic keywords or global geotags, which only marginally support the hyperlocality concept.

Within this paper, we propose a more dynamic and effective approach for hyperlocal content recommendation based on geographic user and content profiling. The idea behind our methodology is based on the criminal investigative methodology of geographic profiling that analyzes the locations of a connected series of crimes in order to identify the likely area where a serial offender resides. In a similar way, location data of a specific user can be analyzed to generate his geographical footprint and geographical media analysis (GMA) can be used to do the same for content items. The user-specific location data can be collected in several ways, e.g., by analyzing GPS activity loggings [3] or social media based events and check-ins [4, 5]. For GMA, both text-based [6] and image-based techniques [7–9] can be used. Due to the generic character of the proposed GEO-profiling architecture, shown in Fig. 1, we do not put any restrictions on the GMA and the geographic user analysis methods. If they are able to produce a geographic content profile (GCP) or geographic user profile (GUP), i.e., a kind of map-based feature layer, they can be integrated in our set-up.

In order to link the GCPs to a particular GUP (and vice versa), our approach makes use of map-based analysis techniques. Firstly, a heat map is created for the GUP and GCPs taking into account several GUP and GCP features, like the trajectory, date and duration of the user activities, and the number of occurrences of a particular place in the media item and its distance to other places in the GCP description. Secondly, we project both heat maps on each other and calculate the heat map score for each of the content items. Finally, the content items with the highest scores are recommended.

The remainder of this paper is organized as follows. Section 2 focuses on the GCP generation and discusses three techniques that we have used to retrieve the geographic content locations. Furthermore, more information is given on the GCP features that are used for the content heat map generation. Subsequently, Sect. 3 proposes our method for GUP generation based on GPS activity logging. As already mentioned before, other techniques for GUP and GCP generation can easily be integrated. Next, Sect. 4 presents the demonstrator of the GEOprofiles-based content recommendations. Finally, Sect. 5 lists the conclusions and points out directions for future work.

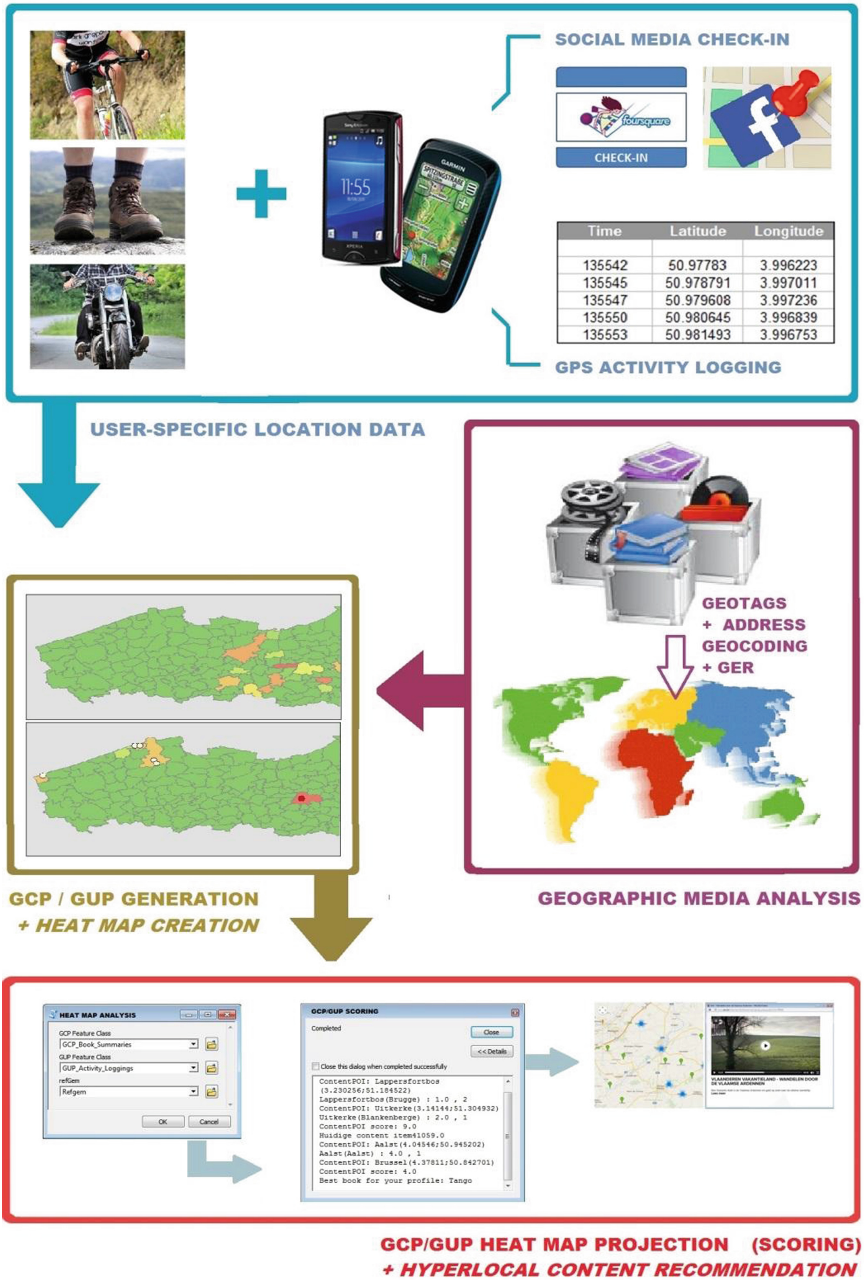

Fig. 1. General architecture of map-based hyperlocal content recommendation using geographic user and content profiling.

2 Geographic Content Profiles

For the generation of the Geographic Content Profiles (GCPs), three different GMA techniques are investigated. Important to remark is that, within this paper, we only focus on textual content. However, also image, video and audio content can be used in the proposed architecture.

First of all, we focus on existing geotagged media. Currently, however, only a small portion of online content is geotagged (mostly pictures and video, and only limited number of text documents and web pages). Furthermore, we observed that most of the geotagged content is only labeled with a geographic tag at the global level, i.e., not within the content itself, which limits its practical applicability in our set-up. Secondly, we focus on address geocoding of available addresses in the content. Some document structuring is needed too easily extract these elements. Finally, we perform geographic named entity recognition on the textual part of the content. This micro-level technique is by far the most valuable one for hyperlocal content recommendation.

2.1 Geotagged Media Content

Over the last decade, geotagging is becoming used more and more. The popularity of the geotag is on the rise by the increased use of Internet capable mobile devices with built-in GPS functionality. The geotag itself is a form of metadata which marks a multimedia object, such as an image, video or text message, with its location information (longitude and latitude coordinates). The majority of recent capture devices are able to automatically assign these kinds of tags.

The huge benefit of geotagged media is that it allows multimedia objects to be browsed and arranged geographically. Photo-sharing websites such as Flickr (www.flickr.com) and Panoramio (www.panoramio.com), for example, provide millions of geotagged images contributed by people from all over the world. In order to retrieve the multimedia data that is related to a specific location, one can choose from several social media web services that support geo-based queries, such as the PANORAMIO geo-picture service and the DBPedia-based FlickrWrappr service [10].

Up till now, the number of geotags for textual documents and web pages is still limited, and the geographic annotation for this kind of content is mostly a manually, error-prone process. However, we expect more and more mechanisms coming soon that will facilitate the textual geotagging process.

A good example of a geotagged webpage is shown in Fig. 2. This webpage of the video archive of Vlaanderen Vakantieland (i.e., a touristic program on Flemish television), has a geotag for each of its episodes. Additional textual information and addresses are linked to each of these geotags, which can be analyzed using the address geocoding and geographic entity recognition that are discussed in the next sections. The set of locations for each episode, i.e. the ‘global’ geotag, the address location(s) and/or the geographic entities detected in the episode’s textual description, will be fed to the GCP generator.

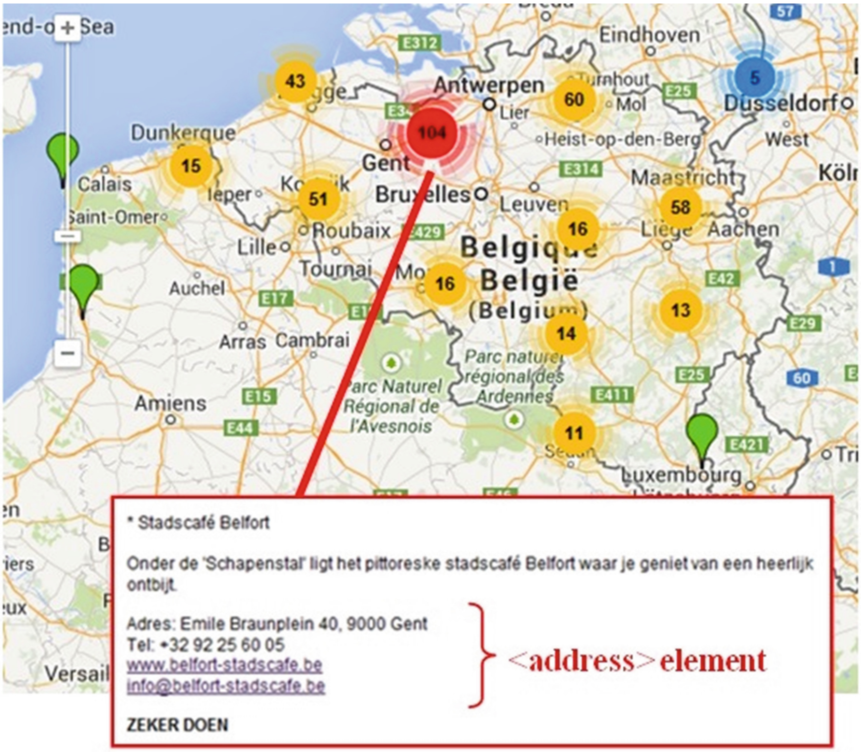

Fig. 2. Geotagged media content – a video archive of a touristic program on Flemish television

2.2 Address Geocoding

Address geocoding is the process of determining an estimated latitude and longitude position for the location of a street address. In the example given in Fig. 2, address geocoding can be used to convert the address element "Emile Braunplein 40, 9000 Gent" into the coordinates (51.053657, 3.723612) of this place.

Address geocoding process. First of all, a parser will break down the address element into a number of components. Then, address standardization identifies each address component (e.g., street number, street name, city and zip code) and places them in order. Finally, the values for each address component are matched to the reference database and an estimate of the spatial location is given. Several matching problems can occur during this process, e.g., misspelled street names, outdated reference data, and incorrect numbers.

In order to automatize the address geocoding process, several address geocoding web services can be used, such as the ArcGis Geocoder[1] and the Mapquest Geocoding

[1] http://geocode.arcgis.com/arcgis/index.html.

API.[2] The former one is integrated in ESRI's geospatial processing programs and can be used, for example, to automatically geocode a table of addresses. Our GCP generator, which is scripted in ArcPy, also makes use of this ArcGis Geocoder to convert address elements in their corresponding locations.

2.3 Geographic Entity Recognition (GER)

Named Entity Recognition (NER) labels sequences of words in a text belonging to predefined categories such as the names of persons, organizations, and locations. NER plays a significant role in many application domains, such as information extraction, summary generation, document classification and internet search optimization. In our work, NER is used for creating a geographical representation of a text or web page, based on the geographic entities that can be detected using NER techniques.

In broad terms, two main types of NER techniques can be distinguished: knowledge based and learning based NER. Knowledge NER techniques use regular expressions, rules and context patterns to detect a particular entity type. In general, these type of NER techniques is very precise and only needs small amount of training data. The drawbacks of Knowledge NER, however, are its expensive development cost and domain dependency. Learning systems, on the other hand, have a higher recall and don't need grammars, but require a lot of training data. For Geographic Entity Recognition (GER), learning based systems have proven to perform best [6, 11, 12].

In our GEOprofiling architecture we use the GER from the iRead + project.[3] This GER is similar to the CLAVIN context-based geotagging service.[4] Both engines extract locations out of structured and unstructured text documents and present geographic features with metadata. The geocoding in both systems is done with help of a gazetteer, i.e. an existing list of entities that automatically can be generated from other data sources. In iRead + , we use an Open Street Map (OSM) gazetteer of Flanders and the discovered geographic entities are weighted according to their occurrence frequency. In Fig. 3, an example is shown of our GER.

2.4 GCP Heat Map Generation

For the generation of the GCP heat maps we make use of a base layer with polygons of interest. In our set-up, these polygons are the town/city boundaries in Flanders. However, other types of polygons, and levels of detail, can be used within the proposed architecture. For each of the GCP locations (geotag, address or geographic entity), we update the polygon that holds it coordinates. In the current set-up, we add the number of occurrences of the GCP location in the content item and weight it with its distance to other places in the GCP description. However, more advanced features/weighting can

[2] http://www.mapquestapi.com/geocoding/.

[3] http://www.iminds.be/en/projects/2014/03/05/iread.

[4] http://clavin.berico.us/clavin-web/.

Fig. 3. Geographic entity recognition within iRead + platform. Location entities are detected using an Open Street Map (OSM) gazetteer of Flanders.

easily be integrated. Finally, to benefit the neighborhood of the GCP locations and to make the content recommendation more flexible, adjacent polygons can be updated with a distance-weighted fraction of the GCP location score.

Figure 4 shows an example of our GCP heat map generation. Two book summaries are projected onto the base map of town/city boundaries in Flanders. The first book summary, of which the GCP locations are represented by the white dots, is geographically spread over multiple polygons. The second book ($\sim$ red dots) its locations are fixed at a single polygon. To calculate the heat map score for each of the content items, the resulting GCP heat map will be projected on the heat map of the Geographic User Profile (GUP), which is discussed in the next section.

3 Geographic User Profiles

The generation of the geographic user profiles (GUP) in our GEOprofiles architecture is based on user-logged activity analysis [13, 14]. Several platforms, like Garmin Connect and RouteYou, provide web services to query the activities of a particular user. Figure 5 shows a data table of one of the author's activities performed during a specific time period. Each of the listed features, such as the startTime, distance and duration, can be used in the weighting of the GUP. Currently, only startTime and duration is used.

The data table in Fig. 5 corresponds to the geographic locations shown on top of Fig. 6. For each of the logged activities, a more detailed view of the trajectory can be retrieved. Such a detailed trajectory can be used to weight all the polygons in the base layer that overlap with the trajectory.

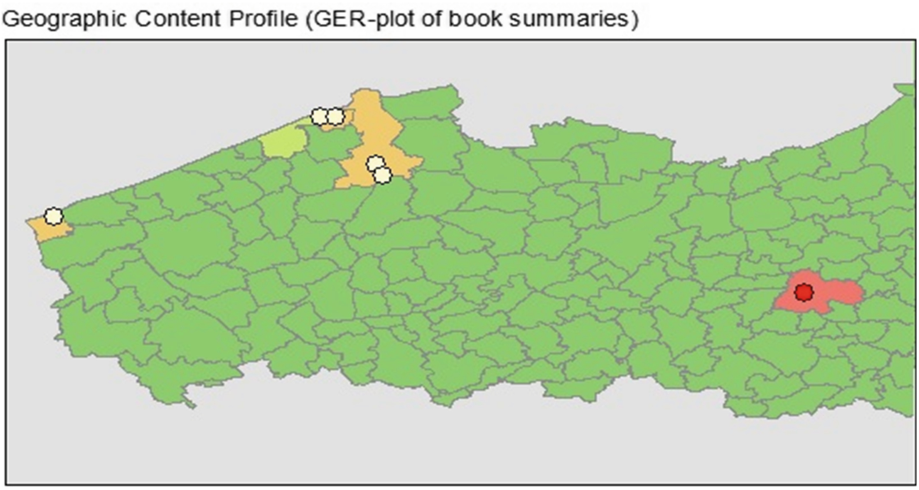

Fig. 4. GCP heat map generation of two book summaries on a base map of town/city boundaries in Flanders.

activityId	startTime	distance	duration	startLat	startLong	calories
253738029	12/15/2012 12:45	45415.94	7670.423	50.97757	3.995826	1589.014
253738039	12/12/2012 13:07	43297.07	8730.193	50.8687	4.174751	1554.885
253738046	12/9/2012 12:12	40000.49	7467.78	51.15179	3.899916	1136.033
253738060	11/25/2012 12:30	40909.83	7280.539	50.9779	3.99608	1351.068
253738078	11/17/2012 9:51	33456.61	5840.738	50.88465	3.892722	1212.882

Fig. 5. Data table of user-specific activity loggings

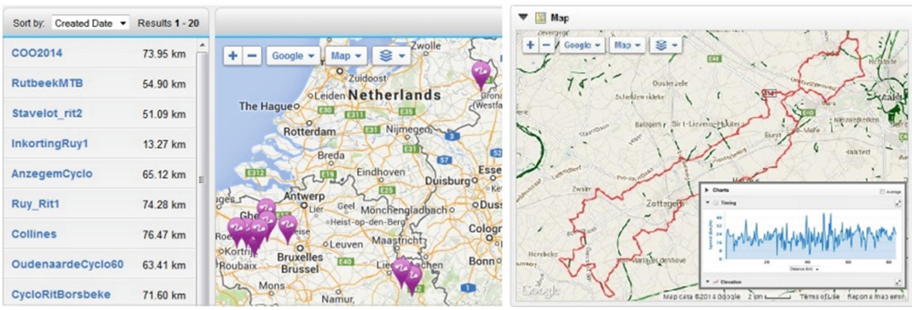

Fig. 6. Detailed view of a specific activity trajectory

An exemplary GUP heat map, on top of the same base map as used in the GCP creation, is shown in Fig. 7. Important to remark is that similar GUP heat maps can be generated for users' social media based events and check-ins. It is even possible to use geographic entity recognition to analyze a user his status messages. Similarly as for

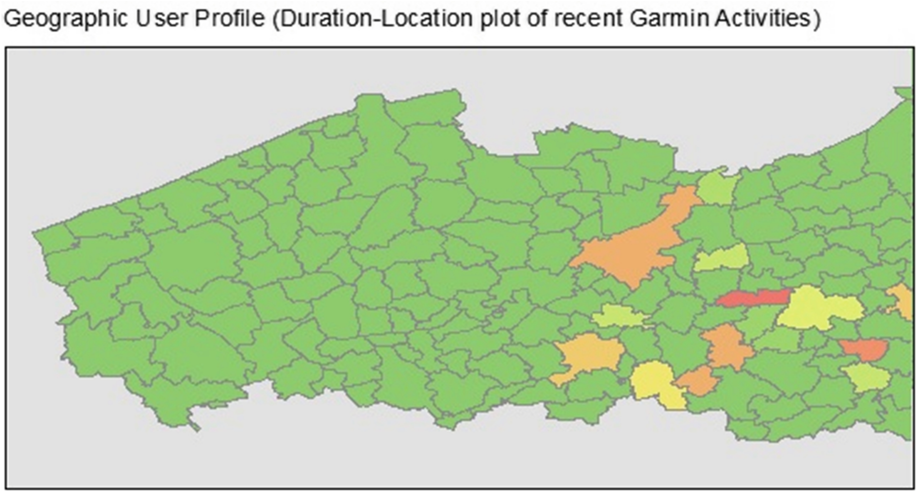

Fig. 7. GUP heat map generation of user-logged activities on a base map of town/city boundaries in Flanders.

the GCP heat maps, adjacent polygons can be updated with a distance-weighted fraction of the GUP location scores.

4 GEOprofiles Demonstrator

In order to evaluate the proposed GEOprofiles architecture for hyperlocal content recommendations, we have developed an ArcPy-ArcGis demonstrator. The demonstrator has been tested on two test cases: the hyperlocal book recommendation (using geographic analysis of book summaries) and hyperlocal recommendation of video episodes from the video archive discussed in Sect. 2.

For both test cases, the heat map score for each content item is calculated by projecting the GCP heat map on the GUP heat map. For each polygon, we multiply the corresponding GCP and GUP polygon values and count them together. Finally, the content items with the highest scores are recommended to the user. The subjective results of both test cases were very positive. Future work will focus on a thorough objective evaluation. Important to remark is that, besides the geographic contextualization of the content, also other contextualization techniques can be added on top of the proposed GEOprofiles architecture.

5 Conclusions

The production and consumption of hyperlocal content is thriving. This paper proposed a novel approach for hyperlocal content recommendation which uses map-based linking of users with content. Key components of the proposed GEOprofiles architecture are the geographic content and user profiles. Several techniques are discussed to extract the geographical locations from the content and user-related data. The GEOprofiling demonstrator, which is evaluated on real activity profiles and two different types of media content, shows the feasibility of the proposed approach.

Acknowledgments. The research activities as described in this paper were funded by Ghent University, iMinds, the Institute for the Promotion of Innovation by Science and Technology in Flanders (IWT), the Fund for Scientific Research-Flanders (FWO-Flanders), the Belgian Federal Science Policy Office, and the EU.

References

1. Paulussen, S., D'heer, E.: Using citizens for community journalism. Journal. Pract. **7**(5), 588–603 (2013)
2. Hu, Y., Farnham, S.D., Monroy-Hernandez, A.: Whoo.ly: facilitating information seeking for hyperlocal communities using social media. In: Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, CHI 2013, pp. 3481–3490. ACM, New York, NY (2013)
3. Biagioni, J., Krumm, J.: Days of our lives: assessing day similarity from location traces. In: Carberry, S., Weibelzahl, S., Micarelli, A., Semeraro, G. (eds.) UMAP 2013. LNCS, vol. 7899, pp. 89–101. Springer, Heidelberg (2013)
4. McKenzie, G., Adams, B., Janowicz, K.: A Thematic Approach to User Similarity Built on Geosocial Check-ins. Lecture Notes in Geoinformation and Cartography - Geographic Information Science at the Heart of Europe, pp. 39–53. Springer International Publishing, Heidelberg (2013)
5. Van Canneyt, S., Van Laere, O., Schockaert, S., Dhoedt, B.: Using social media to find places of interest: a case study. In: Proceedings of the 1st ACM SIGSPATIAL International Workshop on Crowdsourced and Volunteered Geographic Information, GEOCROWD 2012, pp. 2–8 (2012)
6. Martins, B., Chaves, M.S., Silva, M.J.: Assigning geographical scopes to web pages. In: Losada, D.E., Fernández-Luna, J.M. (eds.) ECIR 2005. LNCS, vol. 3408, pp. 564–567. Springer, Heidelberg (2005)
7. Chen, D.M., Baatz, G., Koser, K., Tsai, S.S., Vedantham, R., Pylvanainen, T., Roimela, K., Chen, X., Bach, J., Pollefeys, M., Girod, B., Grzeszczuk, R.: City-scale landmark identification on mobile devices. In: Proceedings of IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2011, pp. 737–744 (2011)
8. Mamei, M., Rosi, A., Zambonelli, F.: Automatic analysis of geotagged photos for intelligent tourist services. In: Proceedings of the 2010 Sixth International Conference on Intelligent Environment, pp. 146–151 (2010)
9. Yeh, T., Tollmar, K., Darrell, T.: Searching the web with mobile images for location recognition. In: Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR 2004, pp. 76–81 (2004)
10. Becker, C., Bizer, C.: Exploring the geospatial semantic Web with DBpedia mobile. J. Web Semant.: Sci. Serv. Agents World Wide Web **7**(4), 278–286 (2009)
11. Mikheev, A., Moens, M., Grover, C.: Named entity recognition without gazetteers. In: Proceedings of the 9th Conference of the European Chapter of the Association for Computational Linguistics, EACL-1999, pp. 1–8 (1999)

12. Silva, M.J., Martins, B., Chaves, M., Afonso, A.P., Cardoso, N.: Adding geographic scopes to Web resources. Comput. Environ. Urban Syst. **30**, 378–399 (2005)
13. Verstockt, S., Slavkovikj, V., De Potter, P., Vandersmissen, B., Slowack, J., Van de Walle, R.: Automatic GEO-MASHUP generation of outdoor activities. In: Proceedings of 11th International Conference on Advances in Mobile Computing & Multimedia, MoMM 2013, pp. 1–4 (2013)
14. Verstockt, S., Slavkovikj, V., De Potter, P., Van de Walle, R.: Collaborative bike sensing for automatic geographic enrichment. IEEE Signal Process. Mag. **31**(5), 101–111 (2014)