A Step toward Personalized Social Geotagging

A. Di Napoli, F. Gasparetti, A. Micarelli, G. Sansonetti
Department of Computer Science and Automation
Artificial Intelligence Laboratory
Roma Tre University
{alfredo.dinapoli,gaspare,micarel,gsansone}@dia.uniroma3.it

ABSTRACT
In this paper, we propose a framework for personalized access to location-aware services. The system relies on the integration of a Geographic Information Retrieval module with a User Modeling component. We describe an application into a specific field, but the platform is easily usable in different contexts. Currently, the framework has been extending to the mobile domain.

Author Keywords
Geotagging, User Modeling, Ubiquitous Computing

ACM Classification Keywords
H.3.5 Online Information Services: Web-Based Services

INTRODUCTION
Geotagging is a practice developed with *Web 2.0*. It is the operation of adding informative metadata to photos, videos, podcasts, websites or RSS feeds. Through these data, it is possible to obtain the geographic identification of the file. Nowadays, world-mapping is being done by powerful tools such as *Google Street View* [1], so Geotagging is a very common practice for both uploads and downloads, thanks to the convergence between 3g network technology and smartphones, which allows users to obtain and supply information in real-time about the places they are passing through.

In our opinion, some factors do not allow users to fully exploit the potential benefits such services are able to offer. First of all, when the user looks for information about the places she is passing through, she is often overwhelmed with tens of references among which she must find the one actually searched for. Such an operation is made even more complex by the necessarily small size of the screen of mobile devices. An other drawback, which can bother some users, is that the proposed references are usually ranked only by a location-based measure, depending only on the distance from where the user is connecting. In some situations, such a way of presentation of references can fail to satisfy the searchers' information needs. Just think, for example, about one user who wishes to buy a color camera and wants to pay as little as possible. This user will not be interested in the stores with a wide range of products, but rather she will look for those sellers who offer the lowest prices. Another possible situation involves one user searching for libraries where it is possible to find a specific kind of books (*e.g.,* Renaissance Painting). The requests of the two users depend just partially on a location-based measure. Both of them, in fact, are willing to cover more distance in order to obtain exactly what they are looking for. A possible solution would be representing the short- and long-term users' needs, in such a manner to fit the suggestions offered by the system.

The goal of the platform we propose here is to provide all users with satisfactory solutions. Specifically, the system we are working on leverages a modeling component which allows the device to provide the user with the requested information, successfully ranked according to her needs.

Moreover, this modeling component allows the system to suggest recommendations without an explicit request of the user. Another feature of the proposed platform is the ease whereby it can be extended to different domains. Later on, we describe an instance in the field of beaches, but the design choices are such that it can be easily extended to different fields.

The system has been implemented through the open-source framework *Django* [2], entirely written in *Python* language. Now the platform runs on desktop, but we are working on its implementation on mobile devices.

The rest of the paper is organized as follows. In the next section, the proposed approach is illustrated. We then discuss the adaptivity of the framework. In the last sections, we review related work and draw our conclusions on what has been done and what is still to be done for the final system.

THE PROPOSED APPROACH
The main idea behind the system is to supply users with a web application in perfect Web 2.0 style, whereby it is possible to share ideas, advises and opinions on itineraries, sites, hotels, beaches, campings, restaurants and much more. The web application we are talking about manages a geographic map freely taggable by registered users. Members can ac-

[1]http://www.google.it/help/maps/streetview

[2]http://www.djangoproject.com

Workshop SRS'10, February 7, 2010, Hong Kong, China.

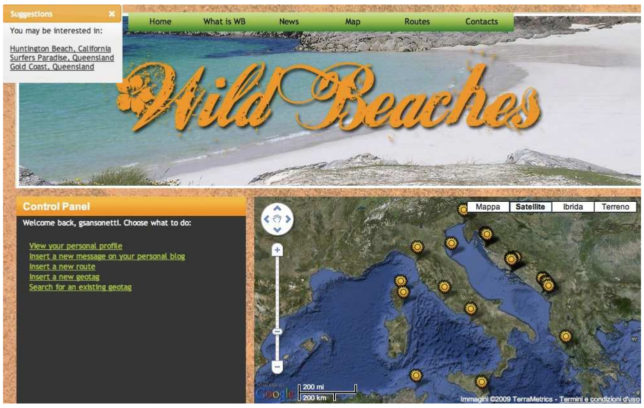

Figure 1. The homepage with the system recommendations.

Figure 2. A Wikipedia description with geographic coordinates.

cess the portal and enter or modify a tag, add descriptions, photos and videos. A discussion forum and a blog complete the community.

As we will see, the design and implementation choices allow the system to be used in different contexts. Figure 1 shows the interface of the instance of the framework in the field of beaches. Tags are denoted by small suns.

Several useful information can be associated to each tag, among which the user who posted the tag, the type of description, the text of the description, the user rating, one or more photos and/or videos.

The *GeoDjango* module of Django has been used to extend the *Object-Relational Mapper* that allows the execution of queries on spatial databases. Moreover, a software model has been created to manage the descriptions posted by users. Such a model allows everyone to choose any point on javascript maps, initially used by *Google Maps*, and to tag that point. This information is saved on the server side as geographic points in the spatial database. The integration with the *GeoIP* module allows to localize the user's location through a correspondence between IP address and geographic coordinates.

In order to populate the platform a Geographic Information Retrieval library has been realized. The implemented methods are the following ones:

1. A Geo-Retrieval module has been structured to satisfy the following workflow:
   a. some keywords concerning the specific field are extracted;
   b. these keywords are given in input to the Google server;
   c. the results of the query are then analyzed and filtered;
   d. each found KML [3] file is submitted to parsing;
   e. each file can contain different geographic sites;
   f. step d. is repeated until all the files have been parsed.
2. The second method for the initial population of the system is to exploit the public web services to already geotagged data. For example, a plugin of the proposed framework allows one to interface *Wikipedia* API. A Wikipedia page (see Fig. 2) contains descriptions of beaches already posted by users. Most pages include geographic coordinates, so it is possible to exploit them to populate the system. Geographic coordinates can be expressed in different formats, among which decimal degrees and degrees:minutes:seconds. The former is useful for KML files and parsing. However, most websites (among which Wikipedia) employ the latter: in that case a conversion is needed.
3. Eventually, a focused crawler has been realized, that is a crawler for downloading the web pages pertaining to the specific topic [5]. This tool allows one to populate the framework, once the field has been defined.

[3] KML is the acronym of Keyhole Markup Language, that is a programming language similar to XML, used by Google to represent geographic sites in Google Earth.

ADAPTIVITY

One distinctive feature of the proposed system is the adaptivity provided by user modeling. Collecting information about individual users is essential for a system to provide the adaptation effect, that is to behave differently for different users. For instance, when the user is interested in strong-wind beaches for surfing, the system can adaptively select and prioritize the most relevant items and recommend them next time she logs in.

Users' interests as well as users' knowledge are the most important features to be modeled in Adaptive Hypermedia systems [2]. In our framework implicit feedback techniques [4] are employed to collect usage data related to user interests. These data are analyzed to build the user profile employed during the recommendation process.

In the chosen domain, the usage data that are possible to collect are the following:

- *Search history:* the set of keywords submitted by the user for querying the search engine;

- *Click history:* the countries of the items chosen by the user;

- *Tag history:* the tags of the items visited or rated by the user;

- *Navigation history:* the items that have been visited or rated by the user.

It is possible to generalize the current domain-dependent data to any kind of retrieval system that provides geospatial metadata, which allows users to navigate through this kind of information. These usage data correspond to the user profile. The system is also able to collect relevance feedback data in the form of item ratings. This is the only feature that allows users to submit explicit feedback in order to be able to learn their needs, interests, and preferences. Aside from requiring additional time during the seeking processes, explicit feedback places a cognitive burden on users and the obtained benefits are not always clear (see, for example, [6]), therefore the effectiveness of explicit techniques may be limited.

Once the data have been collected, they must be included in the recommendation process. The interpretation of these data streams is necessary for the system to tailor itself to each individual's interests. In traditional IR systems, usage data are employed by adjusting the weights of terms in the original query, or by using those data to add words to one query. Relevance feedback is often implemented through the *Rocchio Algorithm* [1].

Recommender systems relying on content-based retrieval systems suffer from the popular vocabulary problem [3]. Vocabulary variety can lead to serious problems during interactions. Different terms may be used to refer to the same object. When users submit queries, they should predict and use the same vocabulary employed when pages have been authored. For this reason, a traditional content-based approach may fail to retrieve relevant information.

Fortunately, in our domain items are also represented by means of tags. Besides traditional geo tags (*e.g.*, geospatial metadata), the user is able to select tags in order to represent features of objects. The vocabulary employed to define the tag set is usually limited and users are often forced to choose one or more tags from this set. Tag history is one of the profile elements that better represent the users' interests. It is collected by analyzing the tags of the items visited or rated by the user, and is also extended to the tags chosen by the user during the search.

Our first approach is to retrieve a subset of relevant items by means of a traditional location-based metric, namely items located closer to the user are judged better. The second step is to rank these items according to a metric that weights how many tags of one item are in common with the user profile. In the current prototype, a *Jaccard similarity coefficient* is used to assess this measure:

$$similarity(TAG_p, TAG_i) = \frac{TAG_p \bigcap TAG_i}{TAG_p \bigcup TAG_i}$$

where TAG_p is the set of tags included in the user profile and TAG_i is the set of items associated with the item i.

The other histories are exploited in order to perform proactive recommendation. Whenever users log in, the system suggests a set of recommendations related to the current interests (see Fig. 1). If the user logs in through a mobile device that is able to send the current geographic coordinates (*e.g.*, latitude and longitude), this information is used to retrieve items in the user neighborhood by means of a traditional location-based metric. The same personalization process previously described is employed to re-rank the current retrieved set according to the user's interests.

The user has the chance to log in via a traditional browser, where data about geographic coordinates are not available or it is not possible to easily transfer these data to the web servers. In this case, the histories related to searches, clicks and past navigations are exploited to retrieve a set if items to be suggested in the graphical UI. In particular, the set of keywords used to query the system in the previous interactions are used to retrieve a first set of results. The matching is based on a IR system with a cosine similarity measure [1]. This result set is then ranked according to the most visited countries and items visited or rated by the user. Basically, the items located in the countries or in the neighborhood of previous items visited by the user are promoted in the list.

This straightforward approach is able to extend the simple location- based approach employed in traditional map services with geotagging support to include elements related to the current users' interests and preferences.

RELATED WORK

Currently, some software applications already allow users to receive information and suggestions concerning the locations from where they connect. *Plazes* [4] is a website where users can realize where other users are and what they are doing. In this way, it is possible to share information and to post next plans. These operations can be done on the website *http://plazes.com*, or through the cell phone. Users can also link to groups to keep in touch with friends and Plazes. The general idea behind this service is to allow everyone to monitor places and people throughout their lives.

The website *http://googlesightseeing.com* is a portal where members can publish information and comments on places of some interest. Specifically, descriptions concern artistic or funny places. Members can also post itineraries or photos by publishing KML files.

Google has just announced that location will be "a first class component of all the products" the company develops. The new Google mobile homepage will shortly report the voice

[4] http://plazes.com

What's Nearby? [5] to allow users to receive personalized search suggestions based on the location. This way, if the user will search for a specific product, the first results will be related to stores where the product is available and nearest to the user's location.

Moreover, the same homepage will include the voice *Near Me Now* [6] to allow everyone to obtain information concerning ATMs, restaurants, bars. The new version of Google mobile maps for *Android* [7] will include this technology.

Another project promoted by Google concerns the 2D mobile barcode format as a camera-friendly tool to find out local businesses and save money in nearby stores [8]. Google has already sent 100.000 mobile barcodes to all American stores. The idea is to show the mobile barcodes in shop windows. Users might be interested in some item, take a picture of the barcode through their cameraphone, and quickly receive info, ratings and coupon.

CONCLUSIONS

In this paper, we have presented a software application which allows personalized access to information available on the web, concerning places and services nearby to the user's current location. In perfect Web 2.0 style, such an application allows everyone to contribute by entering new places or by adding ratings, comments, photos and videos.

The impressions of the first users are satisfactory. Notwithstanding, a deeper evaluation is needed, for which an opportune test-bed must be defined.

The system is still under development. Among other issues, we are extending portal applications to mobile users on *iPhone* and *Android* platforms.

REFERENCES

1. R. A. Baeza-Yates and B. Ribeiro-Neto. *Modern Information Retrieval*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1999.

2. P. Brusilovsky and E. Millán. User models for adaptive hypermedia and adaptive educational systems. In P. Brusilovsky, A. Kobsa, and W. Nejdl, editors, *The Adaptive Web, Methods and Strategies of Web Personalization*, volume 4321 of *Lecture Notes in Computer Science*, pages 3–53. Springer, 2007.

3. L. M. G. G. W. Furnas, T. K. Landauer and S. T. Dumais. The vocabulary problem in human-system communication. *Commun. ACM*, 30(11):964–971, 1987.

4. D. Kelly and J. Teevan. Implicit feedback for inferring user preference: a bibliography. *SIGIR Forum*, 37(2):18–28, 2003.

5. A. Micarelli and F. Gasparetti. Adaptive focused crawling. In P. Brusilovsky, A. Kobsa, and W. Nejdl, editors, *The Adaptive Web, Methods and Strategies of Web Personalization*, volume 4321 of *Lecture Notes in Computer Science*, pages 231–262. Springer, 2007.

6. R. White, J. M. Jose, and I. Ruthven. Comparing explicit and implicit feedback techniques for web retrieval: Trec-10 interactive track report. In *TREC*, 2001.

[5] http://google-latlong.blogspot.com
[6] http://www.betanews.com
[7] http://www.android.com
[8] http://www.intomobile.com