

RecoMap: An Interactive and Adaptive Map-Based Recommender

Andrea Ballatore
University College Dublin
School of Computer Science
and Informatics
Belfield, Dublin 4, Ireland
andrea.ballatore@ucd.ie

Gavin McArdle
University College Dublin
School of Computer Science
and Informatics
Belfield, Dublin 4, Ireland
gavin.mcardle@ucd.ie

Caitriona Kelly
National University of Ireland
Galway
University Road, Galway,
Ireland
c.kelly15@nuigalway.ie

Michela Bertolotto
University College Dublin
School of Computer Science
and Informatics
Belfield, Dublin 4, Ireland
michela.bertolotto@ucd.ie

ABSTRACT

With the growing availability of geo-referenced information on the Web, the problem of spatial information overload has attracted interest both in the commercial and academic world. In order to tackle this issue, personalisation techniques can be used to tailor spatial contents based upon user interests. RecoMap, the system described in this paper, deducts user interests by monitoring user interaction and context to provide personalised spatial recommendations. After an overview of existing recommendation systems within the geospatial domain, the novel approach adopted by RecoMap to produce such recommendations is described. A case study related to a university campus setting is used to outline an application of this technique. Details of the implementation and initial testing of this prototype are provided.

Categories and Subject Descriptors

H.5.2 [User Interfaces (D.2.2, H.1.2, I.3.6)]: User-centered design; I.2.1 [Applications and Expert Systems (H.4, J)]: Cartography; H.5.1 [Multimedia Information Systems]: Hypertext navigation and maps

General Terms

Geographical Information Systems, User profiling, Map personalisation, Recommender systems

Keywords

Geographical Information Systems, GIS, User Profiling, Per-

sonalisation, Recommender systems

1. INTRODUCTION

Due to the steady growth of location-based services, spatial information overload has increasingly become problematic [18]. In the geo-spatial domain this issue is addressed through user profiling techniques aimed at modelling user interests [10]. Ideally, a geographical information system should be able to tailor personalised maps and functionality for a specific user. This is often achieved through the use of non-automated system customisation options. However, this approach requires explicit input, which can be time consuming, distorted by subjectivity and distract the user from their main task [17]. Traditional systems often allow the user to toggle off layers of map content which may be irrelevant to the current task and contributing to display clutter [8]. This level of explicit customisation is of limited use as there is still potential for information overload to be problematic with customisation occurring at the semantic level, rather than on an individual object basis.

By implicitly profiling the user, an insight into their preferences can be gained which permits the automatic personalisation of information and map content. As the personalisation is implicit, there is a larger scope for preference variability. As a result, such personalisation can be performed at the finer-grained object level. Furthermore, as it is non-intrusive, it does not disrupt the current task of the user [6] and as they must interact with the system to obtain the required information, implicit profiling has an intrinsic 100% completion statistic. There are numerous possibilities for spatial systems to gain an implicit insight into a user's information preferences, intentions, and context based on their interactions with map contents.

Several Web-based systems employ an implicit profiling approach based on user actions such as link clicking, page printing, or booking, which act as interest indicators in the contents of a webpage. This approach has also been developed in the geo-spatial domain based on user interactions with map contents [10, 4, 16]. Different indicators have differing strengths and weaknesses depending on the circum-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SAC'10 March 22-26, 2010, Sierre, Switzerland.

Copyright 2010 ACM 978-1-60558-638-0/10/03 ...\$10.00.

stances of their application. This paper describes one such approach called RecoMap. This is a new application which presents personalised and adaptive maps to users, based not only on their preferences but also on their context which are both observed implicitly. In order to demonstrate the benefits of this approach, a prototype, to assist users of an interactive university campus map, has been implemented.

This paper is organised as follows: Section 2 discusses related work in the area of user profiling and spatial recommendation systems. Section 3 outlines the system architecture and implementation with detailed description of the core algorithm used in RecoMap. Section 4 details the prototype which has been developed to evaluate the approach. In Section 5, conclusions on the paper are made and directions for future work are identified and discussed.

2. RELATED WORK

Understanding user's behavior constitutes the core of several interface personalisation techniques. Web user profiling has attracted huge interest both in academia and industry - Kobsa analyses the e-commerce economic boom related to user profiling in the late 90s [7] - and many techniques have been developed and successfully deployed during the last decade [12, 2]. Numerous approaches to filter out irrelevant information from large data sets through user profiling have been proposed. According to Wu et al. [17], two main types of profiling have been identified: *static* and *dynamic*. The former, which aims at collecting characteristics from users themselves, is intrinsically intrusive into the user interaction, prone to subjectivity and fails to reflect the volatile nature of interest [17]. The latter is the process of deducting user interests solely through analysis of interaction with the system. These approaches are usually utilised in specific applications such as personalisation and recommendation.

A consistent body of work has also been carried out in the field of traditional recommendation systems to infer user interests. Balabanovic [3] has defined as *content-based* a recommendation based solely on a profile built up by analyzing the content of items which a specific user has rated in the past, whilst a *collaborative* recommendation relies on a similarity measure among user profiles, without taking the item content into account. Most recommendation systems combine both to overcome the limitations of each approach [14, 3]. Albadvi et al. have conducted an extensive survey on these hybrid techniques [1].

Despite the extensive body of research in web personalisation and recommendation, comparatively little work has been carried out in the field of location-based recommendation, whose commercial potential has become apparent in recent years [15]. To date, the systems which have emerged are diverse and not widely adopted. Hippie [13] is a personalised location-based content delivery service, tailored for museums. It adopts a tagging system to give the user information related to their position. On the other hand, in the field of implicit feedback analysis, CoMPASS [16] is a GIS application that monitors user interaction to recommend semantic groups of objects, such as layers, to users. Similarly, the system described by Mac Aoidh [9] registers user mouse interactions to elicit interests, with the intention of providing recommendations at the object level. PILGRIM [4] is another spatial-aware application which recommends web resources to users based on their location, relying on the spatial dimension of Web documents, with the assumption

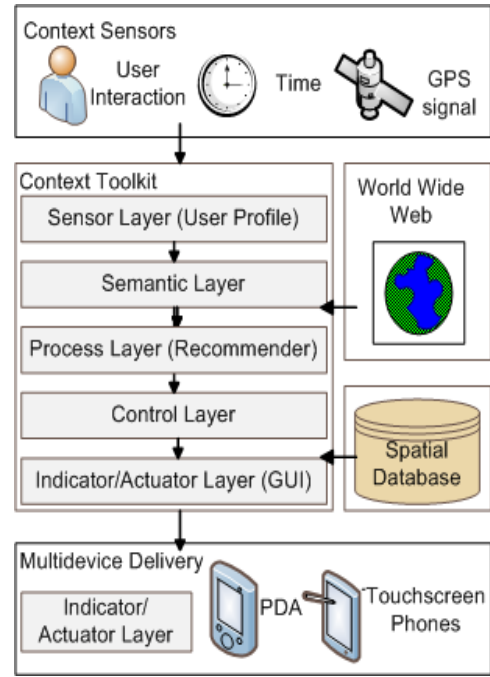


Figure 1: System architecture

that proximity is relevant in interest determination.

Unlike the systems described above, RecoMap combines dynamic and static profiling to provide the user with personalised spatial recommendations through a combination of user context management and user profiling. Additionally, RecoMap extends previous work on implicit feedback analysis [9, 16] by considering context when making personalised recommendations. This approach has received little attention thus far. While Hippie recommends contents based on user location, it does not build a user profile to analyse user behaviour over time, which is a key feature of RecoMap. Similarly, while PILGRIM is focused on web pages, RecoMap also recommends complex spatial items (buildings, parks, points of interests, etc). Furthermore, CoMPASS operates at the semantic level and, like Mac Aoidh's tool, does not take context into account while RecoMap infers user interests from their context at the object level.

3. RECOMAP

RecoMap consists of a modular interactive application that renders a vector map and recommends spatial items and associated information by inferring user interests from implicit feedback indicators and user context. The system architecture, shown in Figure 1, is based on the context management toolkit proposed by Zimmermann et al. [19]. This toolkit integrates user modeling with context modeling and offers a design architecture very suitable for such applications. The *context sensors* collect information and stream it to the *context management system*, whose internal details are described in the following section.

The main component of RecoMap is the **Context Management System (CMS)**, which receives information about the user context from the sensors and provides recommendations to the user. The Context Toolkit is the core of the CMS and the 5 layers which constitute it are described be-

low.

The **Sensor layer** receives information about the user context (user location, an interest radius calculated upon the user speed, time of the day), and their interactions the GUI (mouse clicks and other events). These details are collected in the user profile and are available for the other layers to process.

The semantic structure of the system is defined in the **Semantic layer**. The user world is a set of *items* stored in a spatial database. An item is either a complex geometry (a polygon) or a point of interest representing a geographical entity which the user might have interest in and interact with. Certain items (such as restaurants, shops, etc) can be associated with relevant resources on the World Wide Web (such as web sites or web services). These items are displayed on an interactive digital map (Figure 2).

The **Process Layer** monitors the evolution of the user profile as it is updated by the Sensor Layer. It assigns an interest score to each item located within the current interest radius. The interest score α for the item i is calculated with Equation (1) taking into account both historical interactions (*interaction*) and current user routing distance from the item (*proximity*)

$$\begin{aligned} \alpha_i &= P_i P_R + I_i I_R \\ P_R + I_R &= 1, \quad P_i \in [0, 1], \quad I_i \in [0, 1] \end{aligned} \quad (1)$$

P_R and I_R are respectively the *proximity ratio* and the *interaction ratio*, meaning the weight the system attributes to each indicator. The preliminary evaluation has been conducted with $P_R = 0.2$ and $I_R = 0.8$, emphasising interaction over proximity. Furthermore, the proximity score P_i and the interaction score I_i are normalised between the maximum and minimum score among the items within the proximity score, as illustrated in Equation (2).

$$\begin{aligned} P_i &= \frac{|P_i| - |P_{min}|}{|P_{max}| - |P_{min}|}, \quad |P_i| \in [0, \infty) \\ I_i &= \frac{|I_i| - |I_{min}|}{|I_{max}| - |I_{min}|}, \quad |I_i| \in [0, \infty) \end{aligned} \quad (2)$$

$|I_i|$ is a non-negative real number incremented by a fixed value at every interaction, representing the degree of interest expressed toward the item i (e.g. a mouse click directly on the item increments $|I_i|$ by 0.01 and a click on a recommended item by 0.05). Similarly, $|P_i|$ is intended to represent the ‘spatial interest’ shown implicitly by the user position in the physical world: the more the user spends time close to an item, the higher $|P_i|$ becomes. Given the volatile nature of user interests [17], a time decay function based upon the days elapsed since the last interaction with the item i is applied on $|I_i|$ and $|P_i|$.

When a certain condition occurs in the user context and profile (e.g. user heads toward a new area), the **Control layer** triggers an adaptive action. For example, when the user either starts to explore a new area or alters the user profile by interacting with the map, the control layer obtains new *recommended items*, which have the highest interest score at a given time for a given location.

The **Indicator/Actuator layer** contains the Graphical User Interface (GUI) of RecoMap. The GUI is compatible on multiple platforms including PDAs, tablet PCs and iPhones. As shown in Figure 2, it consists of an interactive map and an information panel. The map has a set of

intuitive functions including zoom, pan and drag. The recommended items are visually highlighted with a bold outline and a stronger colour. The current user location and their historical locations from the Sensor Layer are presented with a trace and a highlighted icon which helps the user visualise the items’ spatial layout.

The information panel, on the right hand side of the GUI, consists of a list of html links and previews of recommended items. When a recommended item is clicked, the panel displays information associated with it and the system acknowledges a successful recommendation by incrementing the interaction score of the item via the Control Layer. Below this, the browser window shows further details of the item the user is interacting with.

4. PROTOTYPE

The prototype has been implemented using OpenMapTM, a JavaTM open-source framework specific for geographical information systems (<http://openmap.bbn.com>). PostGIS, an open-source spatial extension for the PostgreSQL database management system, is used for storing and accessing spatial items and user profiles. The campus dataset used in the prototype has been extracted from OpenStreetMap (<http://www.openstreetmap.org>).

A campus navigation and recommendation tool was developed to demonstrate how RecoMap can be used. The prototype consists of an interactive map of University College Dublin. The vector map data consists of polygons, polylines and points of interest representing the footprints of campus buildings, roads, walkways and various amenities. Additionally, web resources associated with the different on-campus amenities are maintained. A user session starts when the user logs into the system. RecoMap makes recommendations based on user position and on the information previously recorded in the user profile. The recommended items are highlighted in red on the map interface and are also listed in the tabbed panel to the right of the map (Figure 2). The location, time and interactions with the interface are continuously recorded and used to update the profile. As the user location changes, so too do their recommendations. Should the actual user interests change, the time decay function reduces the relevance of previous older interests and more recent ones promptly become more prominent in the profile. When the user logs out of the system, their session is terminated and all the recorded interactions are permanently stored in their profile to be used in future sessions.

5. CONCLUSIONS AND FUTURE WORK

This paper has described RecoMap, a spatial recommendation system that relies on implicit feedback analysis [9, 16] in order to infer user interests and provide spatial recommendations. Its main contribution involves extending current techniques by incorporating user context to develop novel algorithms and personalise spatial contents. Based on a context management framework [19], a layered approach is adopted. RecoMap has separate components for sensing the user context and for recording interaction in order to build a dynamic user profile. The algorithm at the core of the system assigns interest scores to spatial items by combining previous interactions and proximity. A time decay function is also applied to handle shifts in user interests over time.

A prototype has been developed to offer a campus navi-

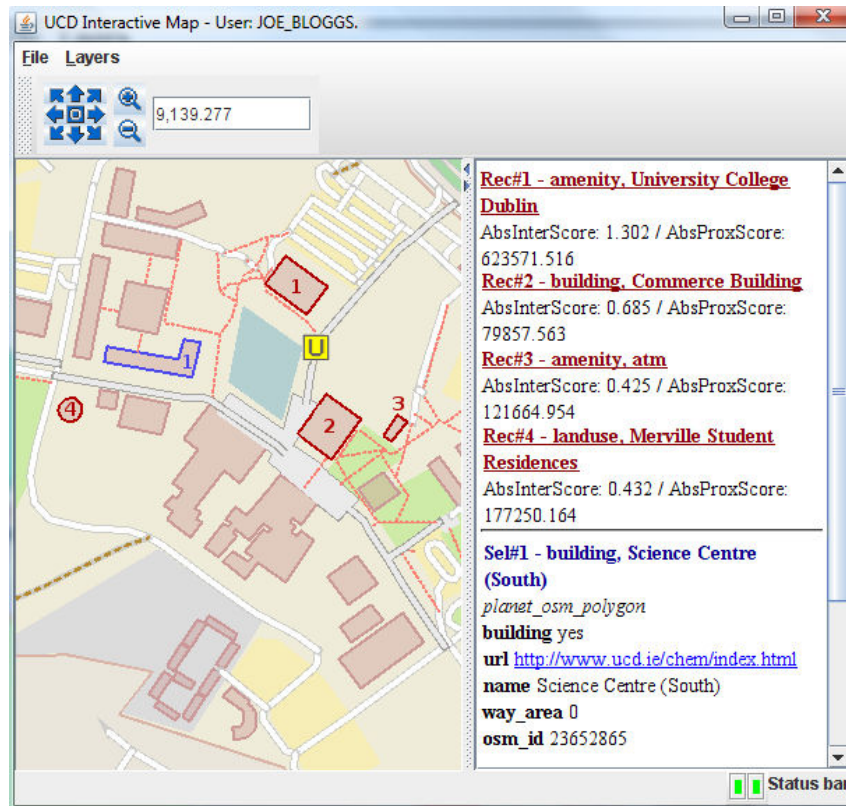


Figure 2: RecoMap: Campus Navigation Example

gation and recommendation tool. It appropriately demonstrates the application of RecoMap in a specific domain by recommending items at the map level to users of an interactive campus map. Preferences are obtained implicitly and combined with context in order to make suitable recommendations. In order to test the prototype, a set of navigational paths representing movements of specific categories of users were recorded on the campus at University College Dublin. These paths were used in conjunction with mouse clicks to simulate several users interacting with the system. Thus, user profiles were automatically generated and used to make recommendations. These simulations have illustrated the potential value and effectiveness of the approach followed by RecoMap.

The next step to evaluate the impact of the RecoMap platform is to conduct a longitudinal study using the campus prototype and test subjects. This study will also enable refactoring, based on user and system feedback, to take place in order to strengthen the algorithm and implementation. Research in the area of user personalisation in Location-Based Services continues to grow. Within this domain the use of ontologies [11] for offering alternative recommendations needs to be investigated. Ontologies will enable, not only previously visited map items to be recommended, but also items that the ontology classes as similar. Additionally, the role of group profiles [5], which could further improve the quality of recommendations, needs to be studied. Determining similarities between individual user profiles enables grouping to take place. Using these collective profiles will increase the relevance of the recommendations and enable

new items to be suggested to users of RecoMap.

6. ACKNOWLEDGMENTS

Research presented in this paper was funded by a Strategic Research Cluster grant (07/SRC/I1168) by Science Foundation Ireland under the National Development Plan. The authors gratefully acknowledge this support.

7. REFERENCES

- [1] A. Albadvi and M. Shahbazi. A hybrid recommendation technique based on product category attributes. *Expert Syst. Appl.*, 36(9):11480–11488, 2009.
- [2] M. Albanese, A. Picariello, C. Sansone, and L. Sansone. Web personalization based on static information and dynamic user behavior. In *WIDM '04: Proceedings of the 6th annual ACM international workshop on Web information and data management*, pages 80–87, New York, NY, USA, 2004. ACM Press.
- [3] M. Balabanovic and Y. Shoham. Fab: Content-based, collaborative recommendation. *Communications of the ACM*, 40:66–72, 1997.
- [4] M. Brunato and R. Battiti. Pilgrim: A location broker and mobility-aware recommendation system. *Pervasive Computing and Communications, IEEE International Conference on*, 0:265, 2003.
- [5] G. Carenini, J. Smith, and D. Poole. Towards more conversational and collaborative recommender systems. In *Proceedings of the 8th international*

- conference on Intelligent user interfaces, pages 12–18. ACM New York, NY, USA, 2003.
- [6] D. Kelly and J. Teevan. Implicit feedback for inferring user preference: a bibliography. In *ACM SIGIR Forum*, volume 37, pages 18–28. ACM New York, NY, USA, 2003.
 - [7] A. Kobsa. Generic user modeling systems. *User Modeling and User-Adapted Interaction*, 11(1-2):49–63, 2001.
 - [8] D. Lanter and R. Essinger. User-centered graphical user interface design for gis. *National Center for Geographic Information & Analysis/NCGIA, Technical Paper*, pages 91–6, 1991.
 - [9] E. Mac Aoidh, M. Bertolotto, and D. Wilson. Understanding geospatial interests by visualizing map interaction behavior. *Information Visualization*, 7(3):275–286, 2008.
 - [10] E. Mac Aoidh, M. Bertolotto, and D. C. Wilson. Analysis of implicit interest indicators for spatial data. In *GIS '07: Proceedings of the 15th annual ACM international symposium on Advances in geographic information systems*, pages 1–4, New York, NY, USA, 2007. ACM.
 - [11] S. Middleton, H. Alani, and D. De Roure. Exploiting synergy between ontologies and recommender systems. *Arxiv preprint cs/0204012*, 2002.
 - [12] S. E. Middleton, N. R. Shadbolt, and D. C. De Roure. Ontological user profiling in recommender systems. *ACM Trans. Inf. Syst.*, 22(1):54–88, 2004.
 - [13] R. Oppermann, M. Specht, and I. Jaceniak. Hippie: A nomadic information system. *Lecture Notes in Computer Science*, pages 330–333, 1999.
 - [14] D. Poo, B. Chng, and J.-M. Goh. A hybrid approach for user profiling. *Hawaii International Conference on System Sciences*, 4:103b, 2003.
 - [15] N. Shiode, C. Li, M. Batty, P. Longley, and D. Maguire. The impact and penetration of location-based services. *Telegeoinformatics: location-based computing and services*, pages 349–366, 2004.
 - [16] J. Weakliam, M. Bertolotto, and D. Wilson. Implicit interaction profiling for recommending spatial content. In *GIS '05: Proceedings of the 13th annual ACM international workshop on Geographic information systems*, pages 285–294, New York, NY, USA, 2005. ACM.
 - [17] D. Wu, D. Zhao, and X. Zhang. An Adaptive User Profile Based on Memory Model. In *Web-Age Information Management, 2008. WAIM'08. The Ninth International Conference on*, pages 461–468, 2008.
 - [18] Y. Yang and C. Claramunt. A hybrid approach for spatial web personalization. *Lecture notes in computer science*, 3833:206, 2005.
 - [19] A. Zimmermann, M. Specht, and A. Lorenz. Personalization and context management. *User Modeling and User-Adapted Interaction*, 15(3):275–302, 2005.