



WeKnowIt

Emerging, Collective Intelligence for Personal,
Organisational and Social Use

FP7-215453

D1.4

User Profiling and Personalisation Tools

Dissemination level	Public
Contractual date of delivery	Month 33, 31-12-2010
Actual date of delivery	21-02-2011
Workpackage	WP7, Case studies
Task	T1.4 User profiling and personalization
Type	Report
Approval Status	Approved
Version	1.4
Number of pages	26
Filename	D1.4_2011-02-21_v1.4_TID
Abstract:	<p>This document summarizes work performed in Task T1.4 devoted to User profiling and personalization, which belongs to WP1 developing the Personal Intelligence layer.</p> <p>The information in this document reflects only the author's views and the European Community is not liable for any use that may be made of the information contained therein. The information in this document is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.</p>



co-funded by the European Union

History

Version	Date	Reason	Author
0.1	12/01/2011	TOC	TID
1.3	02/02/2011	Ready for review	TID
1.4	21/02/2011	After review	TID

Author list

Organization	Name	Contact Information	Reason
TID	Manuel Escriche	mev@tid.es	Author
CERTH	Papadopoulos Symeon	papadop@iti.gr	Reviewer

Executive Summary

This document summarizes work performed in Task T1.4 devoted to User profiling and personalization, which belongs to WP1 developing the Personal Intelligence layer.

WeKnowIt is exploring the concept of Collective Intelligence – a form of intelligence emerging from the cooperation of multiple layers of intelligence using a shared knowledge base. In this context, Personal Intelligence is the layer allowing individual users to provide knowledge and access intelligence using WeKnowIt applications.

The work in task T1.4 has focused on two objectives: (a) recommendation services and (b) personalized tour planning.

The effort aimed at the recommendation service has delivered a prototype adapted to the CSG scenario, where Points of Interests (POIs) are featured by a list of characteristics. User's feedback about POIs allows learning the user profile incrementally. This knowledge enables ordering recommendations of POIs according to closest distance to POIs profiles. The prototype has been integrated in the WeKnowIt system. Additionally a related paper has been published.

Research on personalized tour planning has delivered a conceptual design. According to this, a general touristic city model is personalised by using the user profile in order to plan or recommend personalized routes. This work has crystallized in a proposal for a patent, which is under internal review at Telefónica Investigación y Desarrollo. Papers on this work are written but awaiting final review by TID before their publishing.

Abbreviations and Acronyms

API	Application Programming Interface
CSG	Consumer Social Group
WKI	The WeKnowIt System as seen by the users
ER	Emergency Response
GPS	Global Positioning System
POI	Point of Interest
REST	Representational state transfer
SCC	Sheffield City Council
UK	United Kingdom
TID	Telefónica Investigación y Desarrollo
WKI	WeKnowIt
WP	Work Package

Table of Contents

1. Introduction	6
2. User Profiling and Personalization	7
2.1. Recommendations	8
2.1.1. State of the art.....	8
2.1.2. Implementation	11
2.1.3. Evaluation.....	14
2.1.4. Dissemination.....	15
2.2. Personalized Tour Planning	16
2.2.1. State of the art.....	16
2.2.2. Implementation	18
2.2.3. Patent	19
3. Service and integration	20
4. Conclusions	22
5. References.....	23

List of Figures

Figure 1: Conceptual map	7
Figure 2: Wikipedia POI	11
Figure 3: POI browsing	12
Figure 4: Learning profiles	12
Figure 5: Recommending POIs	12
Figure 7: Server use cases and dependencies.....	13
Figure 8: Get Recommendation	13
Figure 9 : Learn Feedback.....	14
Figure 10: Conceptual design	18
Figure 11: Service Architecture	20

1. Introduction

This deliverable describes the output from Task T1.4 devoted to User Profiling and Personalization, which is included in the work package WP1- Personal intelligence, described in section B.1.3.6 – Work package descriptions of Annex I – “Description of Work”.

WeKnowIt is exploring Collective Intelligence as a construct of multiple layers of intelligence cooperating as a single whole. Personal Intelligence is one such layer and is primarily concerned with allowing individual users to provide knowledge and access intelligence using WeKnowIt applications.

In particular, this deliverable describes the approach and implementation of the recommender system that provides personalized recommendations according to the user profile obtained from the user feedback. This approach is also applied to groups. It has been applied to the Consumer Social Group case study.

This deliverable also describes the personalized tour planning concept implemented and submitted as proposal for a patent at TID patent office.

The document has been structured as follows: Section 2 titled “User Profiling and Personalization” gives context by providing a conceptual map, presenting the state of the art on recommendations, and presenting the approach and implementation for WeKnowIt. It follows Section 3 with a brief explanation about the service and its integration. Finally, there are section 4 about conclusions and section 5 with references.

2. User Profiling and Personalization

User modelling is a cross-discipline research field based on a large body of academic and industry work including artificial intelligence, analytical models, and data mining, with the objective of modelling users' preferences and behaviour in their context.

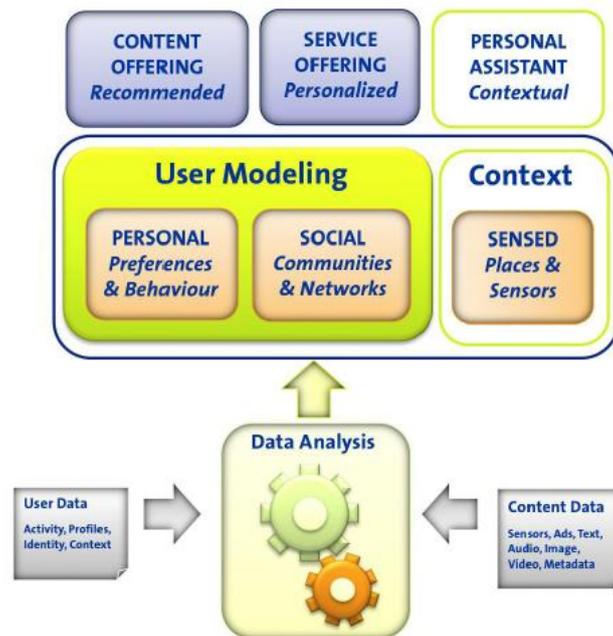


Figure 1: Conceptual map

Some major trends are currently under research, and are already seen as industry trials: the user as part of different communities (social networks: friends, work, people with similar interests or likings), **personalization** of **services** (based on inferred usage patterns) and **content** (preferences, recommendations, collaborative filtering), and context-aware services (sensed urban context, location aware, contextual assistants.)

Personalization is a hot topic, both in research and as an industry trend. The provision of a service personalized to the user's preferences and expected behaviour is a factor that improves user experience and therefore promotes loyalty to the user provider.

This deliverable develops the following topics: recommendation of touristic points of interests and evaluation of recommendation algorithms.

2.1. Recommendations

Recommender systems suggest items based on users' preferences and interests. An early description of recommender systems was "...[a system where] ... people provide recommendations as inputs, which the system then aggregates and directs to appropriate recipients" [37].

The definition of recommender systems has grown broader since these early days. Current recommender systems are more automated and create a user profile in order to propose a small selection of items out of a large variety of options. This profile can be based on a combination of implicit data, i.e. according to the user's patterns of use [2] [50] or, explicit data, where the user briefly, and throughout usage, specifies their preferences to the system [7] [28] [21].

Recommender systems can also differ by the extent to which they engage in a dialog with a user. In "single-shot" recommender systems, each user request is treated independently of previous ones. "Conversational" systems on the other hand are more interactive [10] [31], and users elaborate their requirements over the course of an extended dialog. In particular, the user can supply feedback on the recommended items which influences the next set of recommendations. A discussion of different feedback styles can be found in [22] [42].

These personalized recommender systems have become valuable tools for shifting through large amounts of data. Criteria of retrieving "individualized" as well as "interesting and useful" content have become particularly important [10]. In other words, it is as important that recommendations are interesting, as it is for them to be accurate. If this is done well, the system may gain increased usage and loyalty [32] [23].

Application domains for recommender systems are widely distributed between domains. Previous systems exist in domains such as movies [48], music (e.g. Pandora.com, last.fm), books (whichbook.net, library-thing.com), electronic program guides [27] [2], digital cameras [35], computers [33] and holidays [24].

2.1.1. State of the art

Modern recommender systems use a number of methods. This chapter discusses the following techniques: collaborative filters [32], item-based filters [15], knowledge-based [21] [10], utility-based filters [24] [34] and demographic-based filters [2]. Content-based and collaborative-based filters are the most common types of recommendation algorithms. This is because they are based on rating data, which is relatively easy to collect, and for which there are already established data-sets.

Many recommender systems combine different algorithms in hybrid systems [2] [28] [48] to counterbalance the weaknesses of the individual methods.

Collaborative filters

A recommender system may use correlations between users as a basis for forming predicted ratings of recommended items. That is, recommendations for unseen items are based on information known about ratings of other users, and how these ratings correlate with the ratings of the current user. This approach is called collaborative filtering [48]. Collaborative filtering can be either heuristic-based (also called memory-based) or model-based [1] [9]. Heuristic-based algorithms make predictions based on all the ratings to build a model, such as a naive Bayesian classifier, which is subsequently used to make rating predictions.

Item-based filters

Item-based recommender systems base recommendations on user ratings and similarity between items. That is, while collaborative filters are based on correlations between users, content-based filters are based on correlations between items. Content-based filtering algorithms can be sub-divided into heuristic-based and model-based.

Items are commonly represented as sets of terms or keywords that can be given relative importance using weights. Similarly to collaborative filtering, content-based filtering may use Pearson's correlation [36] or cosine distance [39] to measure the similarity between items based on their keywords. For cosine similarity, the vectors describe the (weighted) frequency of the keywords, or terms in it. Each term defines a direction, or rather a dimension, in the vector space. Similarity is then a measure of proximity between these vectors. For Pearson correlation, the similarity of two items is computed as a weighted sum of all the keywords. Naturally, other types of similarity between items are possible. For example, similarity may be semantically defined, or based on other techniques such as image similarity [18].

Knowledge-based and utility-based filters

Both knowledge-based and utility-based filters can be seen as particular types of content-based filters because item properties (such as price) are used to make recommendations. They are however worthy of their own section since they make a more explicit connection between user requirements and available options than content-based filtering.

Both knowledge- and utility-based filters attempt to make a matching between a user's need and the available options. Their difference is that while utility-based systems require users to do their own mapping between their needs and features of products, knowledge-based systems have deeper domain knowledge. In particular they are able to reason about how a particular item meets a particular user need.

Demographic-based filters

Demographic-based filters (also called stereotype-based filters) use known user characteristics in order to classify users and model user

preferences into classes. A recommender system can define a typical user according to their socio-demographic information, such as postcode, age or gender. Similarly, a combination of demographics can be used to define a stereotype. This stereotype assumes an interest profile, such as particular genres in the movie domain. The user's interests are subsequently classified with regard to how strongly they fit the stereotype, or rather how well they fit several stereotypes.

One could say that demographic-based filters work similarly to collaborative-filters in that they derive similarities between users, but use different types of data. Here similarity is based on demographic information of users rather than their rating patterns.

Hybrids

Many recommender systems use a combination, or hybrid, of methods to counter-balance the weaknesses in each algorithm. Next section discusses the trade-offs involved with each algorithm and return to how hybrid solutions can be used to make best use of respective strengths.

The most common type of hybrid is collaborative/content [4] [25] [27] [40]. As previously mentioned, this is largely due to the lack of availability of rating data. Other solutions have combined demographic user classes and content-based filters using implicit behaviour and explicit preferences [2], collaborative filtering and demographic [28] or collaborative filtering and knowledge-based filters [49].

Burke (2002) [10] discusses different types of hybrids used, and discusses how algorithms can be combined. Some of the methods are order independent, while others give different recommendations depending on which algorithm is applied first. For example, weighted methods are not sensitive to order. In weighted hybrids the scores (or votes) of several recommendation techniques are given relative importance (a weight) and combined together to produce a single recommendation. Other hybrids, such as cascades are sensitive to order. In these hybrids, one recommender refines the recommendations given by another.

2.1.2. Implementation

Based on current state of the art technologies, a recommendation service for Points of Interest has been implemented for the purpose of the CSG scenario. This implementation aims at exploring the idea of using the categories of Point of Interest as main information for user profiling.

The services provided by the recommendation server are used by the CSG mobile phone application, which takes advantage of other WKI services to display information about points of interest, which are drawn from Wikipedia; see Figure 6.

Point of Interests

A Point of Interest¹ defined at Wikipedia is a specific point location that someone may find useful or interesting.

The WeKnowIt system takes advantage from the information about POIs² available at Wikipedia by offering them by mean of WKI services to other services and the applications and tools developed in the CSG scenario.

POIs categories, see **Error! Reference source not found.**, are used by the recommendation server as the mean to profile the POI.



Figure 2: Wikipedia POI

Browsing

The CSG mobile phone application is able to display and present information about POIs from different menu items. In this tool, users can see the location of the POI in a map, can read a textual summary, and display a set of pictures, which enables them to form an idea about the POI.

Besides, the users can express their affection towards the POI in the mobile phone application and can continue exploring or browsing other POIs; see **Error! Reference source not found.**

¹ http://en.wikipedia.org/wiki/Point_of_interest

² Example of POI: http://en.wikipedia.org/wiki/Bank_of_Spain

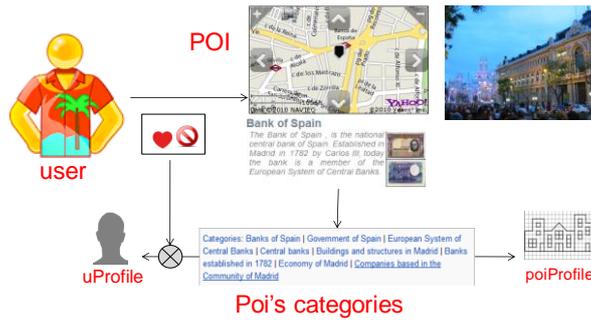


Figure 3: POI browsing

Learning

The recommendation server is able to learn, see **Error! Reference source not found.**, the user profile incrementally by “mixing” the POI profile and the affective feedback from the user. This learning process depends on the list of categories obtained from Wikipedia through WIKI services.

The learning process is implemented by using the ostensive model [11].

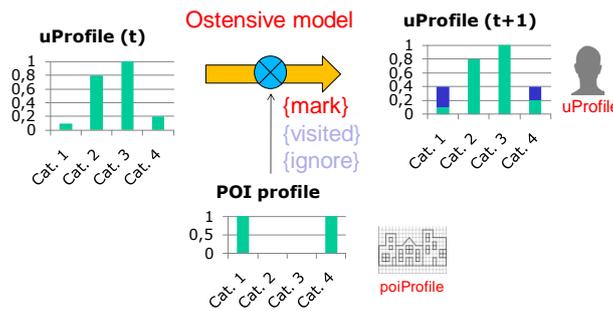


Figure 4: Learning profiles

Recommending

Personalised recommendations are offered to the user under request. The server gets the user profile compared with POI profiles by using the vector space model [39]. This operation determines their similarity indicator, which is used to create an ordered rank of POIs. The list is offered to the user as recommendations, see **Error! Reference source not found.**

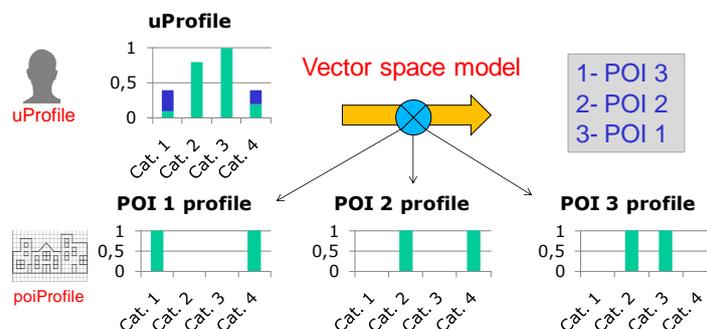


Figure 5: Recommending POIs

Server's use cases and dependencies

The Personalization sever has two main UML use cases, see Figure 6. The use case "get Recommendation" allows the user to get a list of recommendations from the server; and the use case "learn Feedback" allows the user giving feedback about POIs; this last use case gets the list of characteristics by requesting them to the services provided by WP3 – Mass intelligence. The CSG mobile application also gets details about POIs from the services provided by WP3, see Figure 6.

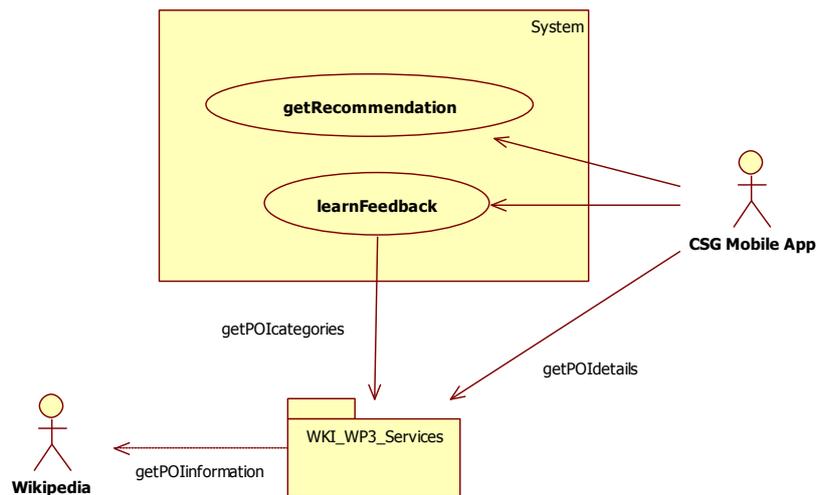


Figure 6: Server use cases and dependencies

The use case "Get Recommendations", see Figure 7, is activated by the user from a menu item. When the server receives the request, it retrieves the user profile from its knowledge base, and continues comparing the user profile with POI profiles. This comparison produces a similarity index, which is used to create an ordered list of POIs as recommendation.

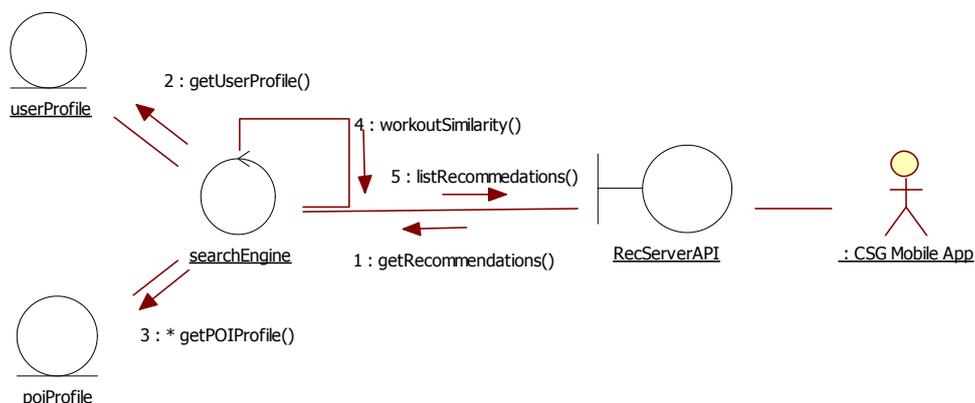


Figure 7: Get Recommendation

The use case "Learn Feedback", see Figure 8, is activated when the user gives feedback about POIs. Upon receiving the request, the server activates the learning engine, which retrieves the POI profile and updates

the user profile by applying algorithms from the ostensive model. The POI profile is previously updated with categories from Wikipedia.

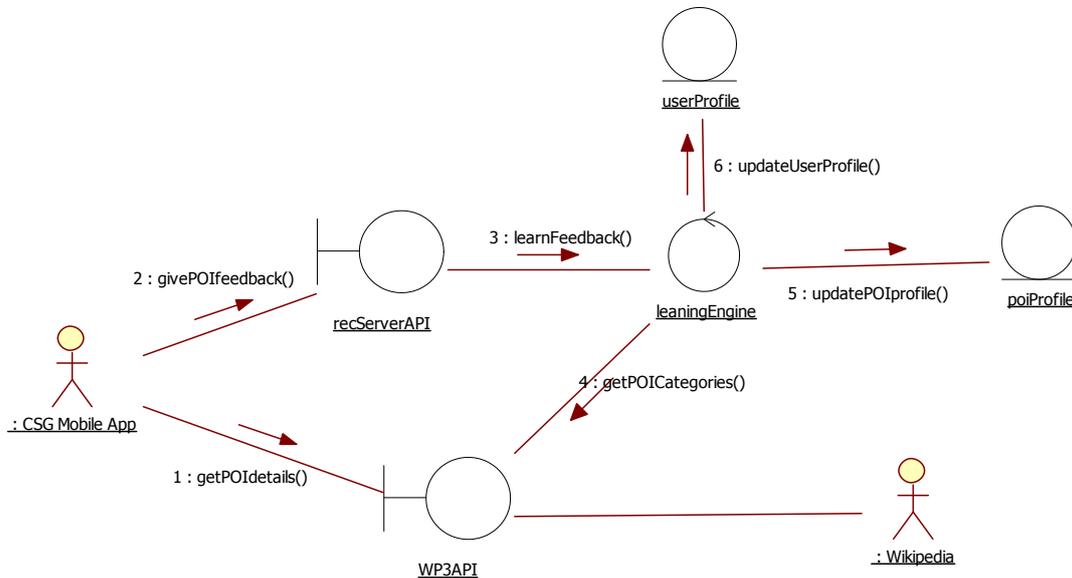


Figure 8 : Learn Feedback

Technology

The recommendation server has been implemented by using following technology components.

Operative System	Red Hat Enterprise Linux 5
Web Server	Tomcat 6
Data Base	MySQL (through Jena SDB storage subsystem)
REST Framework	JAX-RS (Jersey 1.1) + JDOM
Java Platform	Java 6
Semantic Framework	Jena 2.5

The reason to use a semantic framework is the flexibility featured for modelling the knowledge base.

2.1.3. Evaluation

According to deliverable D7.6.2 – “Final consumer and emergency evaluation protocols”, a questionnaire will gather feedback from users about their experience with the mobile phone application in the CSG scenario.

Some questions will be issued on personalized recommendations so as to explore their perception on these suggestions produced by the services.

At internal level, user and group profiles will be extracted, analysed and contrasted with their experience gathered in the questionnaires.

2.1.4. Dissemination

A paper has been published³ with title "Rate it again: Increasing Recommendation accuracy by user re-rating" by the following authors Xavier Amatriain, Josep M. Pujol, Nava Tintarev and Nuria Oliver

Paper's abstract:

"A common approach to designing Recommender Systems (RS) consists of asking users to explicitly rate items in order to collect feedback about their preferences. However, users have been shown to be inconsistent and to introduce a non-negligible amount of natural noise in their ratings that affects the accuracy of the predictions. In this paper, we present a novel approach to improve RS accuracy by reducing the natural noise in the input data via a pre-processing step. In order to quantitatively understand the impact of natural noise, we first analyze the response of common recommendation algorithms to this noise. Next, we propose a novel algorithm to de-noise existing datasets by means of re-rating: i.e. by asking users to rate previously rated items again. This de-noising step yields very significant accuracy improvements. However, re-rating all items in the original dataset is unpractical. Therefore, we study the accuracy gains obtained when re-rating only some of the ratings. In particular, we propose two partial de-noising strategies: data and user-dependent de-noising. Finally, we compare the value of adding a rating of an unseen item vs. re-rating an item. We conclude with a proposal for RS to improve the quality of their user data and hence their accuracy: asking users to re-rate items might, in some circumstances, be more beneficial than asking users to rate unseen items."

This paper comes from laboratory work on datasets and algorithms.

³ RecSys'09, October 23-25, 2009, New York, USA. Copyright 2009 ACM 978-1-60558-435-5/09/10

2.2. Personalized Tour Planning

The tourist domain is one of the domains influenced due to the popularization of the Web. Consumers value the possibility of accessing and comparing options. This allows them to form their personalized travel plans. When planning a city tour, users will perform three different tasks: first, select the places to visit; factor in constraints such as time or budget or contextual factors as weather or transportation; define a plan where all relevant known factors match while allowing for some serendipity.

2.2.1. State of the art

There is a wide range of work related to the theme. We decompose them into four sub-categories: (1) the broader context: urban informatics and activity recognition from movement and web data, (2) previous work that focuses on POI discovery and summarisation using data from the web, (3) analyses of tourist dynamics to generate tourist itineraries, (4) the emerging field of mobile-based location recommender systems.

Mobility tracking and activity recognition

The main input to these systems is users' location, which tends to provide insight into the usage and characteristics of built⁴ environments; for example, Girardin et al. [16] were able to classify locations as tourists or work areas based on aggregated mobile phone usage. Mobile recommender systems, however, are not fully detached from their web-based equivalents. Semantic descriptions and geographic 'hotspots' have been identified based on how web users are annotating geo-tagged Flickr photographs [43]. Peeble et al. [29] develop a framework for learning user-behaviour from crowd-sourced sensor data. Imperfect labels are regrouped via similarity measures. The framework was evaluated by deploying a set of users who logged their activities via audio tags.

Geo-spatial analysis and POI discovery from Web Data

The development of Web 2.0 services has allowed the contribution from user-generated content. These data available on the web is rich with annotations (metadata) describing location (latitude/longitude) and activities (user-input tags).

The metadata can be used to analyze the physical structure of urban environment. Crandal et al. [14] examined the collective photographing behavior of the world's tourists to identify points of interests and rank the locations. These datasets can be used to deduce trip-related information [30] and augment image search results [19].

The user generated content can be analyzed to identify the structures such as the social network between photographers [17] and the collective behavior of web users who tag content [20].

⁴ To differentiate from natural environments

WeKnowIt project has also explored and contributed to this field in Deliverables D3.3 – Mass Classification and Clustering and D3.4 – Report on tools and methods for mass evolution analysis.

Tourist Dynamics and Itinerary Generation

Once POIs have been discovered, paths between them can be computed, which facilitate the exploration of urban environments by tourists. System such as the one proposed by Choudhury et al. [13] aim to automatically generate travel itineraries that “reflect the wisdom of the touring crowds.” To do so, the authors delve further into the data to not only identify POIs, but also estimate visit times and infer paths and transit times between them. The authors propose a pipeline of multiple heuristics that first extracts tourists’ trip data from photos to generate a POI graph, which is used to generate intra-city tours. The GUIDE system [12], on the other hand, operates on pre-defined POIs and generates recommendations that are tailored to the locations that are being recommended. More generally, these systems are requested to be context-aware. In the case of the CyberGuide system [26], context-awareness meant incorporating users’ current and historical location.

Personalisation and Location/Mobile Recommendation

This theme is compiling a growing body of literature. In practice, there are two kinds of systems overlapping. The first kind recommends “while on the move”: these are mobile interfaces to, for example, movie recommender systems [26]. The second kind recommends “locations”, used to find nearby restaurants [38], shops [41], museum exhibits [8] or leisure activities [6].

In the context of the GeoLife (<http://research.microsoft.com/en-us/projects/geolife/>) project, Zheng et al. have built and tested activity and location mobile recommender systems for urban spaces [46]. They built recommendations by reasoning on GPS traces of people’s movements, which exploit correlations between locations or activities.

There are a variety of other works that have addressed the topic of personalization in tourist guides, including [44] and [3]. In general the literature highlights shortcomings on three points: (a) there is no fully comprehensive system, that incorporates POI discovery with city and user modelling to compute personalized itineraries, (b) few systems are easily made general to any location, since they rely on pre-existing POI databases or manual user input, and (c) there is an overriding lack of quantitative and measurable evaluation of the quality of the generated itineraries.

2.2.2. Implementation

The main contribution is a generic, flexible and automated approach to generating personalized tourist routes; a novel approach to construct city models by clustering tagged user photos and discovering POI's; and a way to filter POI's according to a user profile to construct personalized city models.

This has been implemented laboratory work. The conceptual system is composed by three main modules, which are described below.

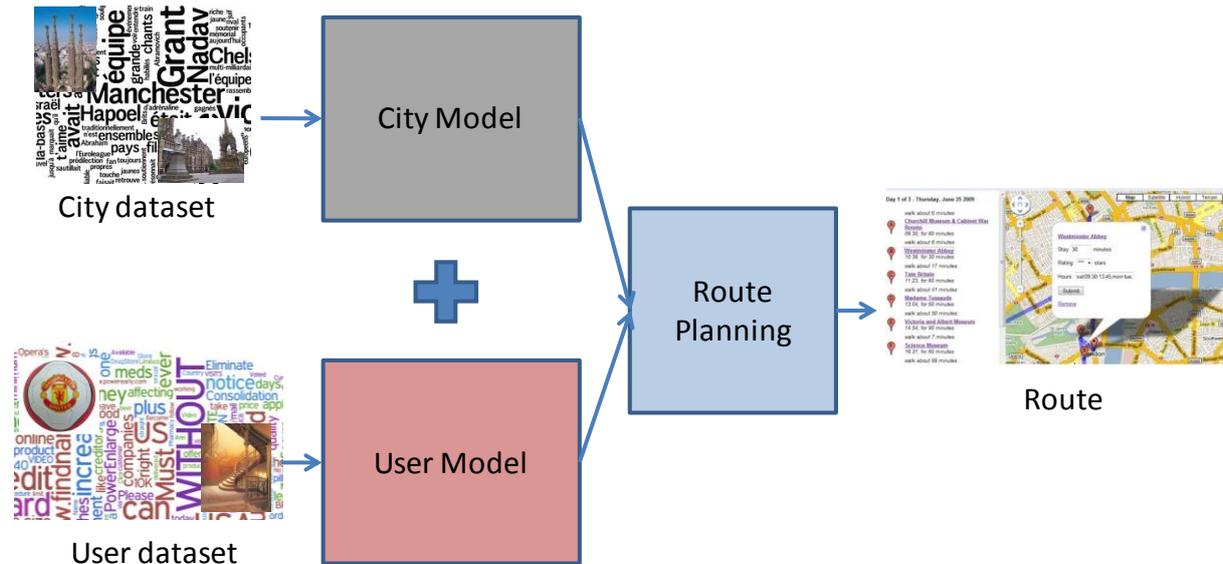


Figure 9: Conceptual design

City Model: the extraction of the tourist city model is made from datasets of geo-located tagged images. The city model is obtained through a clustering process, whose output is a general tourist model of the city. The cluster algorithm takes in two control parameters to define what is considered a relevant POI: a) the minimum number of users agreeing on a given location, and b) the minimum radius for relevant POIs.

This model is composed of Points of Interests, together with their location and associated tags with frequencies. Datasets are filtered to leave out input not related to touristic interests.

User Model: the user data set is converted into a user interest's model. Their interests are extracted from tags and frequencies in her pictures. In this case there is no need of geo-located pictures. The non-existence of given tags in the user model is used to filter out non-relevant POI's from the city model. Thus the generic city model is replaced for a personalized city model with only POI's relevant to the target user.

Route Planning: Once the two previous steps are achieved, personalized touristic routes can be obtained by adding dynamic constraints as weather forecast, opening times, available time, etc to the usual input to the planning process: start and end location, and available time. These constraints are input to the planner, together with the personalized city

model in order to generate the final personalized city tour for the target user.

2.2.3. Patent

The invention described has been submitted by their authors: Xavier Amatriain, Miquel Ramirez and Neal Lathia to TID's patent office. It is under internal review procedure.

- Xavier Amatriain – Telefónica Investigación y Desarrollo - Barcelona
- Miquel Ramirez – Universitat Pompeu Fabra - Barcelona
- Neal Lathia – University College of London

3. Service and integration

The service has been implemented with three main pieces: a learning engine, in charge of updating the user and POI profiles, it uses the ostensive model; a knowledge base, which stores profiles; it is implemented with the semantic web framework Jena; and a recommendation engine, in charge of producing recommendations, by use of the vector space model.

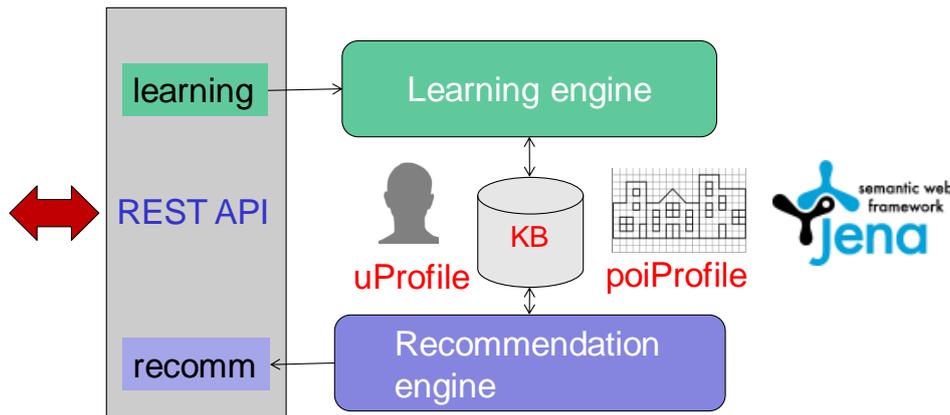


Figure 10: Service Architecture

The service is made available for integration by providing a RESTful⁵ application programming interface (API) within the WeKnowIt architecture.

Next table shows a simplified view of the RESTful API offered by the recommendation server.

Content - POIs
<ul style="list-style-type: none"> • Input <ul style="list-style-type: none"> ○ url:./wp1-ps/content/tourism/{poiId}: <ul style="list-style-type: none"> ▪ POST (create) – PUT (update) ○ Input (category List) • Query <ul style="list-style-type: none"> ○ url:./wp1-ps/content/tourism/{poiId} : GET ○ Output (category List)
Content – Profile
<ul style="list-style-type: none"> • Input <ul style="list-style-type: none"> ○ url:./wp1-ps/user/{userId}/profile/tourism : PUT ○ Input (category List) • Query <ul style="list-style-type: none"> ○ url:./wp1-ps/user/{userId}/profile/tourism : GET ○ Output (category List)
Learning

⁵ REpresentational State Transfer from Roy Fielding's PhD dissertation published in 2000

- Use update
 - [url:./wp1-ps/notify](#) : POST
 - Input (userId, poiId, timestamp, assessment)
 - Assessment (++ , + , -)
- Learning
 - [url:./wp1-ps/profile/learn/{userId}](#) – POST
 - Input (startTime, endTime, useExtProfile, algorithm)
 - Algorithms = constant, exp, linear, log, sigmoid

Recommending

- Query
 - [url:./wp1-ps/user/{userId}/rec/tourism?param](#) : GET
 - Param (number of records)
 - Output (idDocument List)

4. Conclusions

A personalised recommendation server has been provided. Their recommendations are generated by working out items' similarity with the user profile applying the vector space model. The user profile is built by aggregation of weighted POI profiles, where weights express the affection or value feed by the user. This process incrementally learns the user profile by applying the ostensive model.

The service is integrated within the WeKnowIt architecture and made available throughout a RESTful API for the mobile phone application in the Consumer Social Group Scenario.

Laboratory work has been done to achieve a Personalized Tour Planner. As a result a patent has been submitted to TID's patent office by Xavier Amatriain, researcher on Recommender Systems and User Profiles at TID.

5. References

- [1] Adomavicius, G. and Alexander Tuzhilin, A. (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering*, 17:734-749.
- [2] Ardissono, L., Gena, C., Torasso, P., Bellifemine, F., Difino, A., and Negro, B. (2004). *Personalized digital television*, chapter 1, pages 3-26. Kluwer Academic Publishers.
- [3] Ardissono, L., Goy, A., Petrone, G., Segnan, M., Torasso, P. Intrigue: Personalized Recommendation of Tourist Attractions for Desktop and Handset Devices. *Applied Artificial Intelligence*, 2003.
- [4] Balabanovic, M. and Shoham, Y. (1997). Fab: content-based, collaborative recommendation. *Communications of the ACM*, 3:66-72.
- [5] Basu, C., Hirsh, H., and Cohen, W. (1998). Recommendation as classification: Using social and content-based information in recommendation. In Rich, C. and Mostow, J., editors, 15th National Conference on Artificial Intelligence (AAAI-98), page 714-720.
- [6] Bellotti, V., Begole, B., Chi, E.H., and al. Activity-based Serendipitous Recommendations with the Magitti Mobile Leisure Guide. In ACM CHI, Florence, Italy, April 2008.
- [7] Billsus, D. and Pazzani, M.J. (1999). *A personal news agent that talks, learns, and explains*. In Proceedings of the Third International Conference on Autonomous Agents, pages 268-275.
- [8] Bohnert, F. Schmidt, D.F., and Zukerman, I. Spatial processes for recommender systems. In Proceedings of the 21st Joint Conference on Artificial Intelligence, Pasadena, California, July 2009.
- [9] Breese, J. S., Heckerman, D., and Kadie, C. (1998). Empirical analysis of predictive algorithms for collaborative filtering. Pages 43-52.
- [10] Burke, R. (2002). *Hybrid recommender systems: Survey and experiments*. *User Modeling and User-Adapted Interaction*, 12(4):331-370.
- [11] Campbell, I. and Van Rijsbergen, C. J. (1996). The ostensive model of developing information needs. In Proceedings of COLIS-96, 2nd International Conference on Conceptions of Library Science, pages 251-268, Kobenhavn, DK.
- [12] Cheverst, K., Davies, N., Mitchell, K., Friday, A. and Efstratiou, C.. Developing a Context-aware Electronic Tourist Guide: Some issues and experiences. In ACM CHI, The Hague, The Netherlands, 2000.

-
- [13] Choudhury, M.D., Feldman, M., Amer-Yahia, S., Golbandi, N., Lempel, R. and Yu, C.. Automatic Construction of Travel Itineraries using Social Breadcrumbs. In ACM Hypertext, Ontario, Canada, June 2010.
- [14] Crandall, D., Backstrom, L., Huttenlocher, D., and Kleinberg, J. Mapping the World's Photos. In WWW, Madrid, Spain, April 2009.
- [15] Ferman, A. M., Errico, J.H., van Beek, P., and Sezan, M.I. (2002). Content-based filtering and personalization using structured metadata. In JCDL'02: Proceedings of the 2nd ACM/IEEE-CS joint conference on Digital libraries, pages 393-393, New York, NY, USA. ACM Press.
- [16] Fabien Girardin, Josep Blat, Francesco Calabrese, Filippo Dal Fiore, and Carlo Ratti. Digital footprinting: Uncovering tourists with user-generated content. IEEE Pervasive Computing, 7(4):36-43, 2008.
- [17] Golder, S. Measuring Social Networks with Digital Photograph Collections. In ACM Conference on Hypertext and Hypermedia, 2008.
- [18] Jacobs, C.E., Finkelstein, A., and Salesin, D. H. (1995). Fast multiresolution image queuing. In Proceedings of SIGGRAPH 95, pages 277-286.
- [19] Kennedy, L. and Naaman, M. Generating Diverse and Representative Image Search Results for Landmarks. In WWW, Madrid, Spain, April 2008.
- [20] Marlow, C., Naaman, M., Boyd, D. and Davis M.; Position paper, Tagging, Taxonomy, Flickr, Article. In Collaborative Web Tagging Workshop (WWW), 2006.
- [21] McCarthy, K., Reilly, J., McGinty, L., and Smyth, B. (2004). *Thinking positively – explanatory feedback for conversational recommender systems*. In Proceedings of the European Conference on Case-based reasoning (ECCBR-04) Explanation Workshop,, pages 115-124.
- [22] McGinty, L. and Smyth, B. (2002). *Comparison-based recommendation*. Lecture Notes in Computer Science, 2416:575
- [23] McNee, S.M., Lam, S.K., Konstan, J.A., and Riedl, J. (2003). *Interfaces for eliciting new user preferences in recommender systems*. User Modeling, pages pp. 178-187.
- [24] McSherry, D. (2005). Explanation in recommender systems. Artificial Intelligence Review, 24(2):179-197.
- [25] Melville, P., Mooney, R. J., and Nagarajan, R. (2002). Content-boosted collaborative filtering for improved recommendations. In proceedings of the Eighteenth National Conference on Artificial Intelligence (AAAI-2002).
-

-
- [26] Miller, B., Konstan, J.A., and Riedl, J.. PocketLens: Toward a Personal Recommender System. In ACM TOIS, July 2005.
- [27] O'Sullivan, D., Smyth, B., Wilson, D. C., McDonald, K., and Smeaton, A. (2004). *Improving the quality of the personalized electronic program guide*. User Modeling and User-Adapted Interaction, 14:pp. 5-36.
- [28] Pazzani, M. J. (2002). *Commercial applications of machine learning for personalized wireless portals*. In Pacific Rim Conference on Artificial Intelligence.
- [29] Peebles, D. and Lu, H and Lane, N.D. and Choudhury, T. and Campbell, A. T., "Community-guided learning: Exploiting Mobile Sensor Users to Model Human Behavior", in Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence (AAAI-10). 2010.
- [30] Popescu Adrian and Grefenstette Gregory. Deducing trip related information from Flickr. In WWW, Madrid, Spain, April 2009.
- [31] Rafter, R. and Smyth, B. (2005). *Conversational collaborative recommendation – an experimental analysis*. Artif. Intell. Rev, 24(3-4):301-318
- [32] Rashid, A.M., Albert, I., Cosley, D., Lam, S.K., McNee, S.M, Konstan, J.A., and Riedls, J. (2002). *Getting to know you: learning new user preferences in recommender systems*. In International Conference on Intelligent User Interfaces.
- [33] Reilly, J., McCarthy, K., McGinty, L., and Smyth, B. (2004a). Dynamic critiquing. In Funk, P. and Gonzalez-Calero, P. A., editors, ECCBR, volume 3155 of Lecture Notes in Computer Science, pages 763-777. Springer
- [34] Reilly, J., McCarthy, K., McGinty, L., and Smyth, B. (2004b). Explaining compound critiquing. In UK Workshop on Case-based Reasoning (UKCBR-04)
- [35] Reilly, J., McCarthy, K., McGinty, L., and Smyth, B. (2004c). Incremental critiquing. In SGAI International Conference and Innovative Techniques and Applications of Artificial Intelligence
- [36] Resnick, P., Iacovou, N., Suchak, M., Bergstrom, P., and Riedl, J. (1994). *GroupLens: An open architecture for collaborative filtering of netnews*. In Proceedings of ACM 1994 Conference on Computer Supported Cooperative Work.
- [37] Resnick, P. And Varian, H. (1997). *Recommender systems*. Communications of the ACM, 40(3):56-58
- [38] Ricci, F. And Nguyen, Q.N.. Critique-based Mobile Recommender Systems. OGAI Journal, 2005.
-

-
- [39] Salton, G. And McGil, M. (1983). Introduction to Modern Information Retrieval. McGraw-Hill.
 - [40] Symeonidis, P., Nanopoulos, A., and Manolopoulos, Y. (2007). Feature-weighted user model for recommender systems. In User Modeling, pages 97-106.
 - [41] Takeuchi, Y. and Sugimoto, M.. An outdoor Recommendation System based on user location history. In ACM Ubicomp, 2006.
 - [42] Tintarev, N. and Masthoff, J. (2007). *A survey of explanations in recommender systems*. In WPRSIUI associated with ICDE'07
 - [43] Tye Rattenbury, Nathaniel Good, Mor Naaman; Towards automatic extraction of event and place semantics from Flickr tags. SIGIR 2007:103-110
 - [44] Van Setten, M., Pokraev, S. and Koolwaaij, J. Context-aware Recommendations in the Mobile Tourist Application COMPASS. In adaptive Hypermedia and Adaptive Web-based systems, Eindhoven, The Netherlands, 2004.
 - [45] WeKnowIt *Annex I - Description of Work*, ver.1, FP7-215453, 26 Oct. 2007
 - [46] Zheng, Y., Zhang, X., Xie, X. and Ma, W. Mining interesting locations and travel sequences from GPS trajectories. In WWW, Madrid, Spain, April 2008.
 - [47] Zheng, V.W., Zheng, Y., Xie and X., Yang, Q. Collaborative Location and Activity Recommendations with GPS History Data. In WWW, Raleigh, North Carolina, April 2010.
 - [48] Ziegler, C.-N., McNee, S. M., Konstan, J. A., and Lausen, G. (2005). Improving recommendation lists through topic diversification. In WWW'05.
 - [49] Towle, B. and Quinn, C. (2000). Knowledge based recommender systems using explicit user models. In Knowledge-based electronic Markets, Papers from the AAAI Workshop, AAAI Technical Report WS-00-04, pages pp. 74-77, Menlo Park, CA. AAAI Press.
 - [50] Zimmerman, J., Kurapati, K., Buczak, A.L., Schaffer, D., Gutta, S., and Martino, J. (2004). Personalized digital television, chapter 2, pages 27-51. Kluwer Academic Publishers.