PERIODIC REPORT TEMPLATE

PROJECT PERIODIC REPORT

Grant Agreement number: 215453
Project acronym: WeKnowIt
Project title: Emerging, Collective Intelligence for personal, organizational and social use
Funding Scheme: Large Scale Integrating project (IP)
Date of latest version of Annex I against which the assessment will be made: version 2, (03/04/2009)

Periodic report: 1st ☐ 2nd ☐ 3rd ☑ 4th ☐
Period covered: **from** 1st April 2010 (M25) **to** 31th March 2011 (M36)

Name, title and organisation of the scientific representative of the project's coordinator[1]:
Dr. Yiannis Kompatsiaris, Senior Researcher, Informatics and Telematics Institute, Centre for Research & Technology Hellas
Tel: +30 2311 257 774
Fax: +30 2310 474 128
E-mail: ikom@iti.gr
Project website[2] address: http://www.weknowit.eu

[1] Usually the contact person of the coordinator as specified in Art. 8.1. of the grant agreement

[2] The home page of the website should contain the generic European flag and the FP7 logo which are available in electronic format at the Europa website (logo of the European flag: http://europa.eu/abc/symbols/emblem/index_en.htm ; logo of the 7th FP: http://ec.europa.eu/research/fp7/index_en.cfm?pg=logos). The area of activity of the project should also be mentioned.

Table of Contents

1 Publishable summary

1.1 Project synopsis

Exploitation of the knowledge hidden in the user contributed content needs scalable and distributed approaches able to handle the mass amount of available data and generate an optimized 'Intelligence' layer, also called Collective Intelligence. The key purpose of WeKnowIt is to develop novel techniques for exploiting multiple layers of intelligence from user-generated content, which transform the large-scale and poorly structured Social Media to meaningful topics, entities, points of interest, social connections and events. To this end, input from various sources is analysed and combined: from digital content items and contextual information (Media Intelligence), massive user feedback (Mass Intelligence), and users social interaction (Social Intelligence) so as to benefit end-users (Personal Intelligence) and organisations (Organisational Intelligence).

1.2 Summary of activities in the third project year

The project has reached its end with the successful completion of its second implementation phase and the final evaluations phase. Research was directed towards implementing more Collective Intelligence techniques, while a continuous integration plan was carried out for the development of the two use cases prototypes: an Emergency Response application, and a Consumers' Social Group application. Existing Collective Intelligence technologies are integrated into the WeKnowIt system and in the two prototypes, where new functionality features are enabled. The initial versions of the two prototypes have been successfully field tested during the first evaluation cycle, a procedure that has generated more user requirements for the prototypes final versions. The integration of new Collective Intelligence services and the new user requirements have resulted to the finalisation of the two uses cases prototypes. The Emergency Response enables individuals to upload information to the WeKnowIt system about an emergency incident using mobile or desktop devices, allowing intelligent enrichment of the information. This information is then presented to citizens and ER personnel allowing them to understand the incident and make improved real-time decisions and post-incident analysis on the basis of this information. The Consumer Social Group prototype supports travellers in one-day cultural trips by providing desktop and mobile applications in four stages, i.e. an online application during the travel preparation phase, a mobile guidance application during the trip, a landmark recognizer mobile tool used when the traveller is on the go, and finally an online post-travel photo management application. All procedures for testing and final evaluations were efficiently scheduled and the final evaluation phase of the two prototypes was organized and performed. The Emergency Response prototype was evaluated in three rounds: an emergency simulation workshop, a longitudinal user study and citizens evaluation. The evaluations took place in Sheffield and Krakow. The prototype was evaluated by 37 users in total. Concerning the Consumer Social Group case study, the Travel preparation tool has been evaluated in Athens and Krakow from 13 evaluators testing the prototype and answering relevant questionnaires. The Mobile guidance application has been evaluated in Madrid, Krakow and Athens from 13 evaluators, who have participated in 4 hour field tests in the streets of these cities. The WKI image recognizer tool has been evaluated in Krakow, where evaluators tested the efficiency of the tool in recognizing landmarks of Krakow. The Post travel annotation tool is currently undergoing its evaluation. In general, the feedback collected about the concepts implemented for both case studies in research prototypes is positive, though areas of improvement have been identified. During this year the dissemination and exploitation plans were finalized, setting a strong base for harnessing the project's exploitation potential. As a result consortium partners were engaged in the pursuit of market exploitation of the project's new technologies and applications. Dissemination activities included publications, organisation and participation in events, the summer school and contacts

with users and other projects. Exploitation activities included the enhancement of the WeKnowIt User Group, participation in exhibitions (BAPCO, NEM 2010) and industrial events (Mobile World Congress), patent applications, establishment of spin-off companies (Kreuzverweis) and contact with companies in order to establish common exploitation strategies.

An overview of the results according to various aspects and dimensions of WeKnowIt and Collective Intelligence can be seen in Figure 1. The developed Collective Intelligence techniques in most cases exploit links, references and relations among different content items contributed by the users, thus differentiating from the legacy large scale data analysis techniques. Typical examples of such techniques are Flickr-based visual analysis, tag clusters extraction from massive user tagging and social media-based community detection methods. The integration of such different techniques originating from different intelligence layers could potentially leverage Collective Intelligence within diverse usage scenarios.

However, the Collective Intelligence approach proposed in WeKnowIt moves a step further; instead of a mere concatenation of the different layers intelligent methods, it imposes a combination of different intelligence layers within the architecture of some of the developed techniques. Multi-modal analysis is often exploited to enhance the results in each intelligence layer. For example, Mass Intelligence tag clustering results are improved by using Media Intelligence visual analysis features when building graph clusters. As a result the produced clusters are evaluated as more coherent, since they incorporate cross-domain knowledge.

Furthermore, the added value of Collective Intelligence is also evident in the integration level, where the different techniques are combined to produce better results in each case. For instance, evaluation results in WP2 show that geo-tagging through visual and tag analysis yields better localisation results, while Flickr annotations can be improved through automatic localisation and tag recommendations emerging from massive user generated content or the user's social network.

By this Collective Intelligence approach, user generated content items can be analysed by Media Intelligence techniques, resulting in knowledge enhancements about them, as is the case for deriving routes from spatio-temporal analysis using the VIRaL tool. At a higher level, Mass Intelligence techniques are able of analysing massive Web 2.0 users feedback, aiming at extracting implicit knowledge about more generic situations, as is the case with ClustTour and its ability to identify interesting time periods or areas in cities. In overall, the WeKnowIt Collective Intelligence methodology is able to produce enhanced results by:

- exploiting large-scale user contributed content
- combining different layers in building Collective Intelligence techniques
- fusing results from different intelligent layers

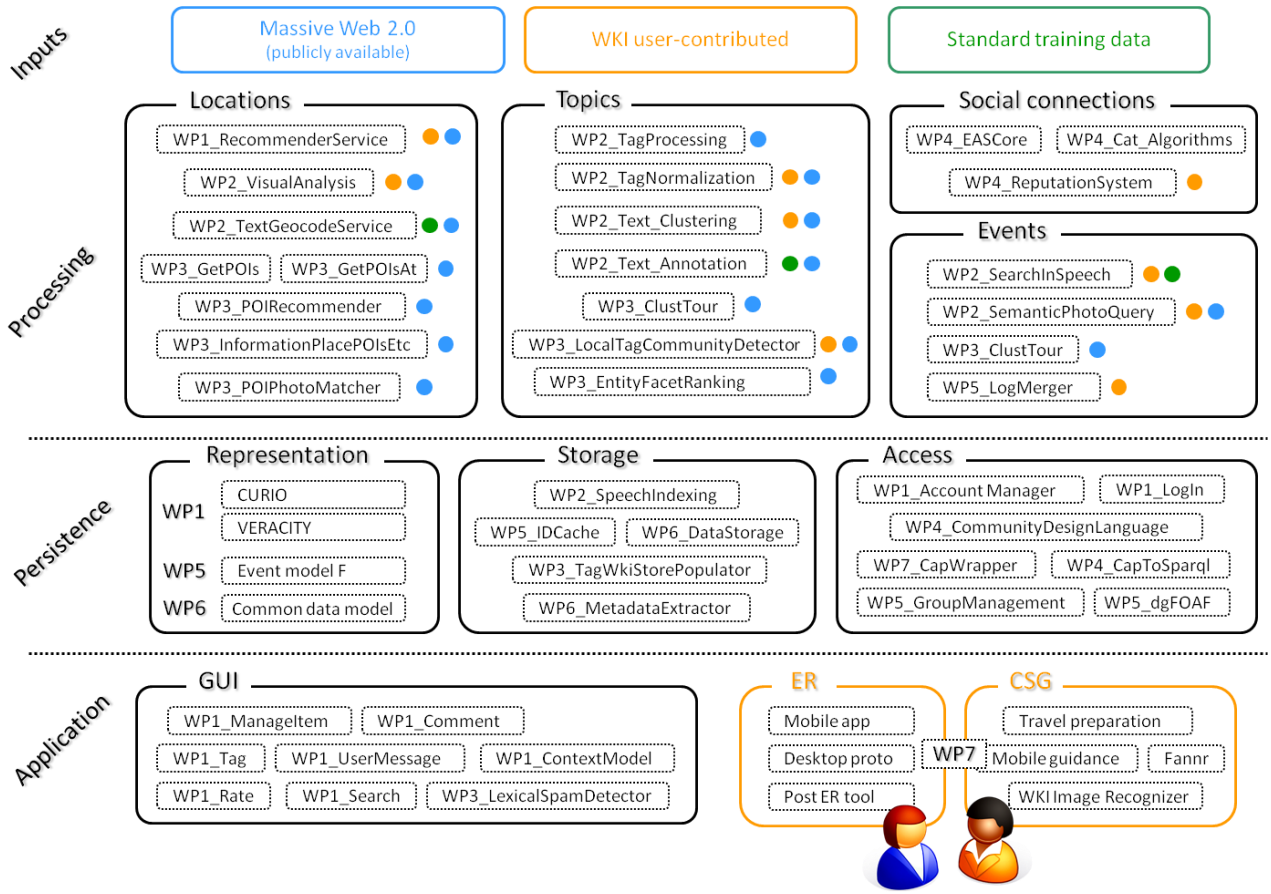

Figure 1 Overview of WeKnowIt services and various aspects.

A web page where available APIs from all partners are listed has been developed (http://mklab.iti.gr/wki-apps/). This page (Fifure 2) describes in detail how the developed WeKnowIt services can be accessed and tested and provides links to responsible partner and person for further information and communication. The objective is to publicly provide to interested users the WeKnowIt services, but also to serve as a "one-stop" site where WeKnowIt specific technologies are presented and can be accessed.

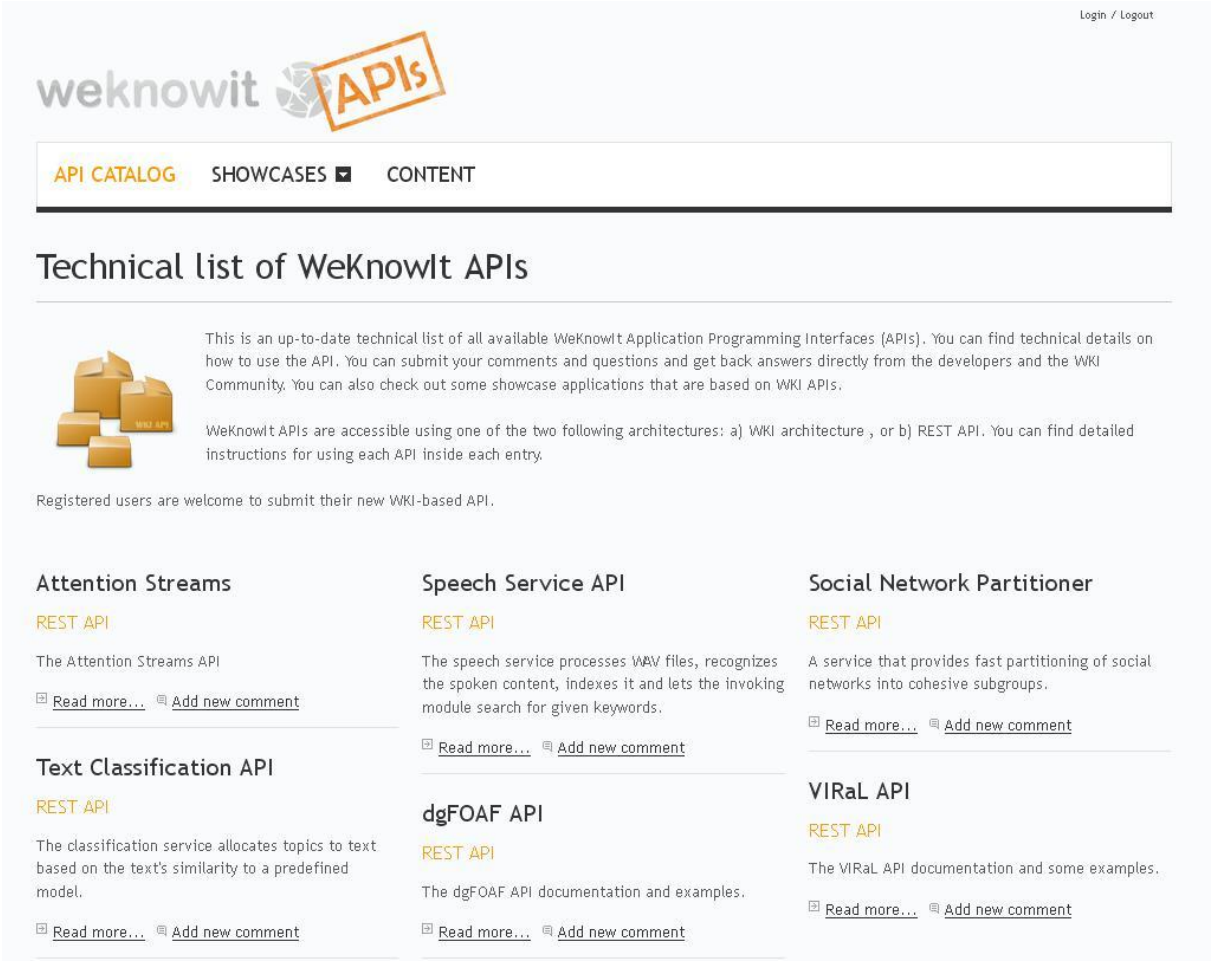

Figure2: The WeKnowIt APIs page.

Personal Intelligence

The Personal Intelligence strand of work has focused on three principal areas. Firstly we have explored how to model users in a real time dynamic environment. We have investigated how to build a multi-scale model of user interests, which can represent both static long-term interests alongside short-term dynamic interests. This allows us to account for both levels when modeling users. We have also investigated how to model the context of the user. Again, context is multi-faceted property of the user and accounts for the users environment, technology and location. We have built a flexible model, which allows WeKnowIt users to specify the context manually (e.g. if the user is in location X they are at the scene of an emergency) or to learn context through the sensor inputs (e.g. given the output of the accelerometer, is the user travelling or staying still). This allows the WeKnowIt system to account for the context of the user and deliver a more tailored user interface. Finally we have also been exploring how to recommend locations, or Points of Interest, to the user on the basis of their interests. Currently we are using Wikipedia to build models of each location and match these models with user interests in order to drive a flexible and relevant set of recommendations.

Media Intelligence

The Media Intelligence layer of WeKnowIt offers a variety of innovative features and tools, including advanced textual and speech content processing, as well as intelligent, automated visual image similarity analysis and retrieval techniques. Textual achievements focus on extraction of

important events from textual documents, by either automatically tagging texts or by indicting texts which can go on to further textual analysis (such as annotation). Speech analysis focuses on quick and efficient interpretation of raw audio and speech content by offering advanced techniques of Out-Of-Vocabulary words identification. Both tasks are proven to be very effective in tackling well-known limitations of traditional tagging and speech recognition systems. Significant efforts have also been spent on visual image retrieval and localization of digital still images through the production of VIRaL application. VIRaL is a photo search engine that can locate your photos on the map, it identifies landmarks and suggests tags to be added to a user photo. VIRaL finds visually similar photos within seconds, without having the user to provide any textual input about the photo. Its main innovation is that it locates any photo on the world map, it automatically identifies any depicted landmark, suggests tags and associates the content to Wikipedia articles. VIRaL's state-of-the-art visual similarity algorithms ensure highly accurate results within a very short time span. In addition the fusion of information is also examined, discussing how to identify where facts concerning a given object (document) are related, i.e. (semantically) similar. The fusion processes which address the combination of reinforcing facts and resolution of conflicting inconsistencies are considered. The fusion of information is examined within the separate media analysis techniques and also the fusion of related facts derived from both textual and visual media. Finally, all developed techniques are evaluated. In order to evaluate the work conducted within the Media Intelligence workpackage, we follow a thorough evaluation process, dealing with all aspects and media involved during the entire project's duration.

Mass Intelligence

A key research activity within Mass Intelligence layer concerned the application of community detection methods that were developed during the first two years of the project on the problem of photo clustering and cluster-based landmark and event recognition. The developed framework has been implemented in ClustTour, an online application demonstrating the results of the framework on a large photo collection of 23 European cities. ClustTour has been extended in order to provide a spatio-temporal content organization layer on top of the photo clusters, and at the same time refine the detected clusters by taking into account the spatio-temporal distribution of photos in the dataset. In that way, ClustTour endows users with enhanced content exploration and browsing capabilities, and at the same time, it improves upon the quality of the presented clusters. The application is online at: www.clusttour.gr. Further achievements in the Mass Intelligence layer includes named entity discovery and ontology-based categorization, event detection in folksonomies and evolution of topics over time in social media, characterization of social media by combining text features with spatial knowledge, event and landmark detection in social media content, user and tag activity analysis over time in social media environments, faceted exploration of image search results, automatic localization of Wikipedia articles, and dynamic analysis of cluster evolution for burst detection.

Social Intelligence

The WeKnowIt Social Intelligence layer covers the processing of the social relations of online communities as well as providing a flexible authorization service for social applications. The analysis of social networks focuses on scalable algorithms for community detection in networks. This capability of inferring social groups from huge communication networks is important for the Emergency Alert Service (EAS) – a prototype application for WeKnowIt's emergency response scenario. The EAS is a location-based service that activates nearby members of the social group of a victim in the case of an emergency. It is a mobile service that works as an emergency call agent, informing social contacts (friends, family, and colleagues) and public authorities about the emergency situation. Current geo-position data and routing information is provided. The service has

been designed privacy-aware and dynamic: Friend lists are calculated automatically with behavioral methods. Friends are only alerted when physically close enough to provide help. The Community Administration Platform is an access control service developed by the Karlsruhe Institute of Technology. In difference to existing approaches, the community administration platform can be plugged into almost any environment as it is realized as external service, allowing defining not only the usual policies and access rights but also the underlying model. A graphical user interface was also developed. Further work in cross-usage of intelligence has resulted to the development of two more Collective Intelligence Services, i.e. the WKI reputation system and a multi-modal recommendation engine. Both these approaches take into account results from other Intelligence Layers services in order to rate or recommend content items or users accordingly.

Organisational Intelligence

In emergency response, the different professional entities such as emergency hotline, police department, fire department, and emergency control center use different systems with their own proprietary data models for events. Using the formal Event-Model-F instead, these systems can commonly represent and effectively communicate event descriptions. The Organisational Intelligence Event-Model-F bases on the foundational ontology DOLCE+DnS UltraLight (DUL) and provides a set of ontology design patterns to represent the different relations of events such as participation of objects in events, mereology, causality, and correlation. It is used in the ER Log merging and management application WERL. It automatically merges different ER logfiles and represents them using the Event-Model-F. Furthermore, semantic information is extracted from their text and a slider-based time filter enables the examination of a particular time interval of the incident. The Event-Model-F is also used in the SemaPlorer++ application, an extension of the SemaPlorer application, for creating and sharing event descriptions. The application allows its users to create event descriptions by clicking on a concept in the ontology on the left hand side, and drag and dropping it on the map. These events can be documented with textual annotations and pictures taken from Flickr.

The WeKnowIt system and use cases

The software framework created for the WeKnowIt System integrates loosely coupled modules produced by each intelligence layer. 40 services are currently deployed and integrated using Enterprise Integration Patterns. Among them, a data storage module is provided, that effectively stores and serves data of various kinds - multimedia files, relational data, Java objects and RDF triples. Within the WeKnowIt System data is exchanged using a Common Data Model, (based on the developed Curio[3] ontology). A custom Java-to-triples mapping mechanism is utilized.

Software solutions implemented within the WeKnowIt System are based on open-source software. Use case applications communicate with the WeKnowIt System via web services REST APIs, that exposes functionality of services integrated within the WeKnowIt System. Additional layers of abstraction of the WeKnowIt System, prepared per use-case, result in easy to use API tailored to the requirements of each application. Features of the publicly available REST API (i.e. its asynchronous nature) facilitate the creation of new end applications. This enabled the rapid creation of additional not initially foreseen applications, which exploit the WeKnowIt technologies, like the WeKnowIt Image Recognizer for mobile phones. The WeKnowIt Image Recognizer is a mobile application that provides the user with detailed information on the location and name of a POI that she photographed. The application uses services developed in WeKnowIt in order to recognize Points-Of-Interest on the picture, determine its geolocation, and determine tags associated with this POI.

[3] http://staffwww.dcs.shef.ac.uk/people/G.Burel/ontologies/curio/

Emergency Response case study and Evaluation

The work on the ER demonstrator has mainly focused on creating a second version of the prototype by taking into account new research functionalities developed throughout the year, and additional requirements and feedback from the users. To respect the multimodal requirement and the need for context-driven interaction two versions of the prototype have been created; a desktop interface and a mobile interface. The mobile interface has been implemented as a customised application for Android Smartphones. The desktop interface has been enriched with new functionalities, such as a messaging service to instantly reach users on their smartphones, the possibility to submit incident reports both from desktop interface and mobile application (with automatic geolocation) and the capability of performing post-incident management searches and analysis. Newly developed intelligent services were integrated such as a tag analyser and a reverse geo-coding web service, services for managing users and groups, as well as services for the background architecture. A new functionality for uploading audio files was added to the mobile application, allowing users to upload messages that can then be listened to from the desktop interface (this functionality is ideal for emergency situations, to allows emergency responders to upload descriptions of the event or comment and to keep the recording of this intelligence for future analysis).

The evaluation of the emergency response application was undertaken in different sessions, following different methodologies and with different participants for each session. Comparing the results of the three evaluation rounds, the final scores are positive especially considering the prototype status of the application. No major usability problems were identified but feedback was given on some aspects off the system that could be improved, and some improvements have already been incorporated in the system. It is very important to notice how the users appreciated the quality of the intelligent suggestions provided by the system, especially for the geolocation and reverse geocoding system and the tag suggestions. During the longitudinal study evaluation was clear how at the beginning most users did not fully understand the reason for tagging or using automatic suggestions, but by the end of the study it had become evident that tagging the information using the automatic tag suggestions was a very quick means to easily retrieve the information at a later stage and communicate valuable information to the persons accessing the system.

Figure 3: ER evaluation exercise.

Consumer Social Group case study

This case study covers a scenario where the focus is on travel or one-day cultural trip events. The scenario is composed of two parts: the first is about preparing the trip with a personal computer at home, and the second part is about supporting the group by means of guidance application running on a mobile phone while having the trip. During the second iteration running along Year 3, the prototype has enhanced its features. The pre-travel tool shows improved usability based on feedback collected in the evaluation stage. Various WeKnowIt services were also integrated, such as, point-of-interest recommendation and hybrid image clustering. Storing of bookmarks was re-implemented using the del.icio.us API, allowing for easier exchange of data between prototypes. In addition, the post-travel prototype, namely Fannr, aids users when annotating their photos after the travel. The mobile guidance application kept improving with enhanced map controls, display of events around the user, access to bookmarks taken during the pre-travel phase, photo cluster visualization and UI polishing. The final version of the CSG prototypes was evaluated by independent users in three different places: Athens, Krakow and Barcelona (Figure 4).

The CSG evaluation process regarding the prototypes indicated that there were no major usability problems identified but the evaluation has pointed out some aspects that can be improved in further development of the prototypes. One of the main conclusions was that the evaluation results depend on the location where the evaluation takes place. In hindsight this is obvious since the prototype is dependent on collective knowledge from social media sites and the support of big cities is much better than for smaller cities. The experience in three cities have been different but helpful to providing a more complete view on how the prototypes can be further improved.

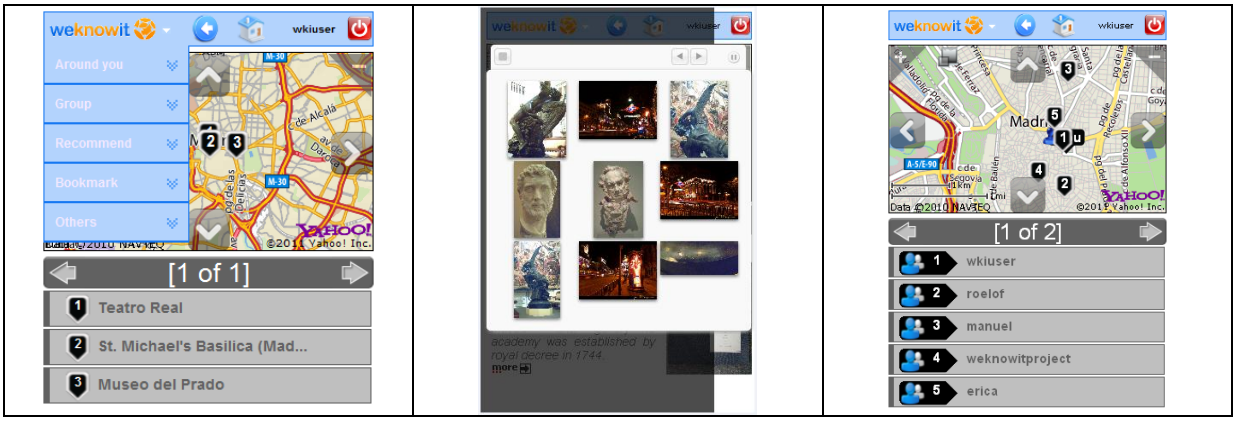

Figure 4: CSG Mobile guidance interface snapshots.

1.3 User Involvement, Promotion and Awareness

During the third year, the WeKnowIt consortium has successfully participated in and organised a number of dissemination events and participated in industrial events in order to promote the project results. A summary is depicted in the table below:

Activity	Number	Notes
Workshops	6	- SETN '10, EVENTS - ACM MM '10, WSM - ICT Event '10, 1 Networking Session - ACM MM '10, EiMM - ACM MM '11, EiMM - ACM MM '11, WSM
Conferences	3	- SMIND WKI conference organization in Warsaw - ACM '10, 2 papers, 2 demos, 1 grand challenge participation - WWW '10, 5 papers
Publications	4	- 29 conference papers - 5 workshop papers - 3 journals - 2 book chapters
Media	4	- Butterfly EUROPE event - Video showcase production - Special event on WeKnowIt, Czech Republic - Video as supplementary multimedia material for the accepted article "Cluster-based Landmark and Event Detection on Tagged Photo Collections" in the IEEE Multimedia magazine.
Summer Schools	1	- SSMS '10
Exhibitions	4	- BAPCO '10 - Infosystem '10 - NEM '10 - MWC '11
Clustering	3	- Common exercises with emergency organizations - Cooperation with Pronto,GLOCAL and VisitoTuscany - Expansion of WeKnowIt User Groupby 9 members

Table 1- Year 3 dissemination and exploitation activities

User testing has been actively pursued during the final evaluation phase of the project. Especially, with the creation of the WeKnowIt User Group, many of its members (such as Bristol City council, Manchester City Council, Niobium Labs, Sboing.net etc.) have engaged in testing the WeKnowIt results, even after the end of the project (e.g. St. Paul's Carnival in Bristol).
Furthermore, users awareness is pursued through the issuing of the monthly WeKnowIt newsletter, presenting all the latest news and events of the project, useful information and tips. Finally the WeKnowIt video showcase has been produced and is publicly released[4].

1.4 Future Work

The WeKnowIt project ended in March 2011 with the implementation of the final prototypes for both use cases and their evaluation and testing. However, testing will be continued after the end of the project, as is the case with the ER prototype (St Paul's Carnival in Bristol). Furthermore, partners are expected to take advantage of the exploitation opportunities that were identified during the project lifecycle. Specifically, bilateral agreements on future synergies between partners of the consortium and external companies are foreseen, since most academic partners of the project have already started exploiting the project outcomes through partnerships with small SMEs (targeted at testing and commercializing individual techniques developed in WeKnowIt). For example, CERTH is discussing with Imagga the licencing of the clustering technology to be used in their StockPondium[5] application. Industrial partners are in the process of embedding WeKnowIt techniques into their products, e.g. Yahoo!'s faceted image search module. They are also targeting to protect their IPRs through patent applications. One spin-off company has been established by University of Koblenz-Landau. Finally, the ClustTour idea was submitted to the OpenFund - a Greek venture capital aiming at financing and supporting promising high-tech start ups - receiving positive feedback regarding the potential for exploiting the underlying technologies. Several more similar activities are carried out by individual partners of the consortium.

1.5 Contacts and Information

For further information visit the project web site http://www.weknowit.eu or send e-mail to the project coordinator Dr.Yiannis Kompatsiaris, ikom@iti.gr.

[4] http://vimeo.com/22334041 , http://vimeo.com/22341074 , http://vimeo.com/22340329
[5] http://www.imagga.com/projects/stockpodium/