

WeKnowIt Annual Report 2010



www.weknowit.eu

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).

Summary of Activities

The project is reaching the end of its second implementation 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. All procedures for testing and final evaluations have been already discussed, in order to efficiently organize the final evaluation phase of the two prototypes. During this year the dissemination and exploitation plans were finalized, setting a strong base for harnessing the project's exploitation potential. In the next year the main focus will be on the finalization of the two uses cases prototypes, their evaluation implementing new field tests, and finally the pursuit of market exploitation of the project's new technologies and applications.

Different Intelligence layers building Collective Intelligence

The WeKnowIt approach towards Collective Intelligence builds on two aspects: mass content availability provided by a lot of users and availability of analysis techniques and results from different layers. Cooperation between different modalities brings real value to the end users. A new Collective Intelligence paradigm is introduced which is able to combine the different intelligent layers, and exploit their interactions and synergies in order to effectively harness Collective Intelligence in the integration level. The WeKnowIt Collective Intelligence techniques exploit links, references and relations among different content items contributed by the users, thus differentiating from the legacy large scale data analysis techniques.

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.

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 (Fig. 1). 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.




Fig. 1: VIRaL application: Retrieval of similar images and localization of query image on Google Maps.

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 is detailed in [1], and ClustTour, an online application demonstrating the results of the framework on a large photo collection of 23 European cities was presented in [2]. For instance, Fig. 2 presents the top 10 landmarks detected by the proposed framework for the city of Barcelona. In addition, Fig. 3 presents several events that were detected for three European cities. The application is online at: www.clusttour.gr and ongoing research is conducted to refine the quality of the obtained results.




Fig. 2: Top ten landmarks detected by ClustTour for Barcelona




Fig. 3: Sample events detected by ClustTour for three European cities

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 (see Fig. 4). 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. More information can be found in [3], [4]. 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 (Fig. 5).




Fig. 4: Emergency Alert Service: Homescreen alert widget on an Android smartphone (left) and emergency notification (right)

Session Control

Elements

Transitive Sets

Relations

Meta Sets

Tests

Access Conditions

Check Access

Command line

Commands sent

Items

Create a new one

SimpleObject ☐

SimpleUser ☐

wkidoc ☐

wkidocgroup ☒

wkiuser ☐

wkiusergroup ☐

Number of items: 6

Delete selected

Properties

Property name	Property value
id	151
externalid	wkidocgroup
subtype	TRANS
type	S

Contained elements:

Contains intransitively

wkidocgroup1
wkidocgroup2

Contains transitively

Contains in deep

wkidocgroup1
wkidocgroup2

Other elements

SimpleAccessRelation
SimpleObject
SimpleTest
SimpleUser
mSimpleUserSimpleAccessR...
o1
o2
u1
u2
wkilac1
wkidoc
wkidoc1
wkidoc2
wkidoc_docgroup
wkidocgroup

Other elements

SimpleAccessRelation
SimpleObject
SimpleTest
SimpleUser
mSimpleUserSimpleAccessR...
o1
o2
u1
u2
wkilac1
wkidoc
wkidoc1
wkidoc2
wkidoc_docgroup
wkidocgroup

Commit

Abort

Request in progress: none

Fig. 5: Graphical User Interface of the Access Control Service. It depicts the transitive set wkidocgroup. This set contains directly two subgroups, wkidocgroup{1|2} and no intransitive assignments.

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 [5] 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 depicted in Fig. 6. 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.



Fig. 6: Snapshot of the WERL front-end main screen

The Event-Model-F is also used in the SemaPlorer++ application, an extension of the SemaPlorer [6] application, for creating and sharing event descriptions. The application is shown in Fig. 7 and 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.




Fig. 7: Screenshot of the SemaPlorer++ Application for Creating and Sharing Emergency Response Event Descriptions

WeKnowIt system integration and infrastructure technology

The software framework created for the WeKnowIt System integrates loosely coupled modules produced by each intelligence layer (Fig. 8). More than 30 services are currently deployed and integrated using Enterprise Integration Patterns. Among them, a data storage

module (Fig. 9) 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 Curio¹ ontology). A custom Java-to-triples mapping mechanism is utilized. Software solutions implemented within the WeKnowIt System are based on open-source software.




Fig. 8: The WKI System architecture - with use-case applications




Fig. 9: The WeKnowIt Data Storage module

¹ <http://staffwww.dcs.shef.ac.uk/people/G.Burel/ontologies/curio/>



Fig. 10: WKI Image Recognizer - a mobile application recognizing landmarks and Points of Interest

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 (Fig. 10). The WeKnowIt Image Recognizer is a mobile applications 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.

Emergency Response Demonstrator and field tests

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 (see Fig. 11 and 12) 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, see Fig. 13) and the capability of performing post-incident management searches and analysis (see Fig. 14).




Figure 11: The ER desktop interface




Figure 12: Visualising an uploaded geolocated picture

Incident Details

Title *

Type of incident Severe weather_Storm

Hazards ☐ yes ☐ no

Hazards details

Casualties ☐ yes ☐ no ☐ not sure

Number of casualties

Notes

Fig. 13: Incident report form

Current Template

Other Templates

New_Incident_Report

hasIncident_Details

Incident_Details

hasType_of_incident

hasHazards

hasTitle

hasinternal_contacts_present

hasHazards_details

hasCasualties

hasNumber_of_casualties

hasNotes

hasexternal_contacts_present

hasIncident_Location

Incident_Location

hasLocation

hasAccess

hasEgress

hasRVP

New Incident Report(3)

Incident Report(11)

Criteria : has Title = fire

Document	has Title
1290980462794	Store on fire
1288887446312	Fire
1288872271069	Some Fire

1290980462794 X

NEW INCIDENT REPORT

Incident_Details

Title

Type of incident

Hazards

Hazards details

Store on fire

Fire_Other

yes

no

Fire may spread to nearby buildings. Risk of explosion - needs to be evaluated by Fire specialists

Fig. 14: Search over incident reports

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 WeKnowIt project recently participated in a simulation of an emergency exercise, in cooperation with the Sheffield City Council and Sheffield Emergency Services. The simulation was of a chemical attack on a major building in the city centre. The emergency services were testing how effectively they could set up decontamination tents and process as many members of the public and sufficiently decontaminate them.

WeKnowIt participated in two capacities. First of all, a member of the team observed the entire incident, from the beginning to dealing with decontaminated people in the Sheffield United football ground. The purpose of this observation was to collect a large set of geo-located images for use in future experiments and to test how well the WeKnowIt ER demonstrator could display and process.

In addition to this five citizen participants were asked to arrive at the incident at a random time and to capture some images using a stripped down version of the WeKnowIt interface to examine how a typical user annotates emergency information in the field.



Fig. 15: WeKnowIt Emergency Response field test



Fig. 16: WeKnowIt Emergency Response exercise

CSG Demonstrator and field tests

The Consumer Social Group case study covers scenarios where a group of people do one-day cultural trip. 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.

The first prototype with the basic features reported last year, undertook two evaluations: one performed by a group of four experts and a second by a group of 22 users contributing to the project as a work session in Barcelona. The outcome gathered is shown in Fig. 17.





Figure 17: 1st phase evaluation of the Mobile Guidance application

The evaluation shows positive assessment in some dimensions but it also points out that further work is needed to complete the application.

The travel preparation prototype was evaluated using two groups of users: a group of 21 WeKnowIt consortium members and a group of 15 users external to the WeKnowIt project. The subjects were asked to fill in a standard usability questionnaire as well as answering questions about specific aspects of the system and open questions for general feedback. During the second iteration running along 2010, the prototype has increased its features. Actually the following features were added in the pre-travel tool: 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. The storing of bookmarks was re-implemented using the del.icio.us API, allowing for easier exchange of data between prototypes. In addition, work was started on a post-travel prototype where users are aided when annotating their photos after the travel. The mobile guidance application kept improving: enhanced map controls, display of events around you, access to bookmarks took in the pre-travel phase, upload of tagged photos to Flickr server, photo cluster visualization and UI polishing. Here is a group of snapshots from the tool (Figs. 18-21):



Second iteration prototype will be evaluated in March 2011 by independent users in three different places: Athens, Krakow and Barcelona.

User Involvement, Promotion and Awareness

This year we have promoted the progress of WeKnowIt in 13 events, a number of which we have organised or sponsored, as listed below:

- Co-organization of **EVENTS 2010: 1st International Workshop on recognising and tracking events on the Web and in real life**, hosted by the 6th Hellenic Conference on Artificial Intelligence conference (SETN2010), 4 May, Athens, GR.
- Sponsorship of the **Summer School on Multimedia Semantic 2010**, 30 Aug.-3 Sept., Amsterdam, NL
- Sponsorship of the **2nd ACM International Workshop on Events in Multimedia**, 25-29 Oct., Firenze, IT

We have also pursued some significant opportunities for gaining visibility in a market-oriented audience, like our participation in one-to-one meetings in the **2010 GSMA Mobile World Congress** industrial event, the representation of the project in the **CeBIT 2010** exhibition, as well as in the **BAPCO 2010** conference and exhibition and in **NEM 2010** summit exhibition.

WeKnowIt was also successfully presented in the worldwide premier multimedia conference, **ACM Multimedia Conference 2010** in Florence with support for 2 workshops (on events

and social media), 2 oral papers (acceptance rate ~16%), 2 demo presentations and 1 Multimedia Grand Challenge participation.

In order that we seek the co-operation with organisations which show interest in following our activities and possibly creating synergies or participating in our evaluation tests, we have started a **WeKnowIt User Group** which – at the time of writing – lists 15 companies.

Interested readers can get more details on our success stories through our website (<http://www.weknowit.eu/>) and even stay tuned to our latest news by subscription to our newsletter.

Future Work – Exploitation Prospects

The following year starts with the final evaluation of the WeKnowIt prototypes. The phase will include evaluation tests with external users of the platform, especially in terms of collaboration with members of the WeKnowIt User Group.

In terms of the project exploitation potential, WeKnowIt will pursue participation in industrial events and contacts in order to establish more synergies. Also 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). 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, namely Kreuzverweis.com, which kicked off last September. 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.

Further Information

The consortium

1. Centre for Research and Technology Hellas (Greece)
2. University of Koblenz-Landau (Germany)
3. The University of Sheffield (United Kingdom)
4. University of Karlsruhe (Germany)
5. Vodafone Panafon - Hellenic Telecommunications Company S.A. (Greece)
6. Software Mind S.A. (Poland)
7. The Sheffield City Council (United Kingdom)
8. Brno University of Technology (Czech Republic)
9. Telefonica I+D (Spain)
10. Yahoo! (Spain)

Contacts and Information

For further information visit the project web site www.weknowit.eu or send email to the project coordinator Dr. Yiannis Kompatsiaris ikom@iti.gr.

Web links

You can visit some of the applications that are part of the WeKnowIt prototype:

The Emergency Response demonstrator: <http://nebula.dcs.shef.ac.uk/weknowit/>

The desktop application of the Consumer Social Group case study:

<http://weknowit.research.yahoo.com/csg>

The mobile CSG application: <http://weknowit.tid.es/WeKnowIt>

The post-travel desktop CSG application: <http://weknowit.research.yahoo.com/fannr/>
WeKnowIt Image Recognizer: <http://www.weknowit.eu/wkiimagerecognizer/>
ClustTour tool: <http://www.clusttour.gr/>
ViRaL tool: <http://www.image.ntua.gr/iva/tools/viral>
SemaPlorer: <http://btc.isweb.uni-koblenz.de/>
STEVE: <http://isweb.uni-koblenz.de/Research/systeme/csxPOI>
ViRaL tool: <http://www.image.ntua.gr/iva/tools/viral>
ER Log Merging: <http://mklab.itl.gr:8080/LogMergerDemo/>

References

- [1] S. Papadopoulos, C. Zigkolis, Y. Kompatsiaris, A. Vakali. "Cluster-based Landmark and Event Detection on Tagged Photo Collections". In IEEE Multimedia Magazine (accepted for publication, October 2010)
- [2] S. Papadopoulos, C. Zigkolis, S. Kapisris, Y. Kompatsiaris, A. Vakali. "ClustTour: city exploration by use of hybrid photo clustering". In Proceedings of MM '10, the international Conference on Multimedia (Firenze, Italy, October 25 - 29, 2010), ACM, New York, NY, 1617-1620
- [3] A. Geyer-Schulz, M. Ovelönne, und A. C. Sonnenbichler, "Getting Help In A Crowd - A Social Emergency Alert Service", in Proceedings of the International Conference on e-Business 2010, p. 207-218, 2010.
- [4] M. Ovelgönne, A. Sonnenbichler, und A. Geyer-Schulz, "Social Emergency Alert Service - A Location-Based Privacy-Aware Personal Safety Service", in Next Generation Mobile Applications, Services and Technologies (NGMAST), 2010 Fourth International Conference on, p. 84-89, 2010.
- [5] Scherp, A, Franz, T, Saathoff, C, Staab, S (2009). F-A Model of Events based on the Foundational Ontology DOLCE+ Ultra Light. K-CAP.
- [6] Schenk, S, Saathoff, C, Staab, S, Scherp, A (2009). SemaPlorer - Interactive semantic exploration of data and media based on a federated cloud infrastructure. J. Web Sem., 7(4).