

Emerging, Collective Intelligence for personal, organisational and social use

Sotiris Diplaris¹, Andreas Sonnenbichler², Tomasz Kaczanowski³, Phivos Mylonas⁵, Ansgar Scherp⁴, Maciej Janik⁴, Symeon Papadopoulos¹, Michael Ovelgönne², Yiannis Kompatsiaris¹

¹ Centre for Research and Technology Hellas – Informatics and Telematics Institute, Thessaloniki, Greece

² Karlsruhe Institute of Technology, Germany

³ Software Mind S.A., Krakow, Poland

⁴ University of Koblenz-Landau, Germany

⁵ National Technical University of Athens, Greece

1. Introduction

Due to advances in communications, mobile devices and Web technologies, it is nowadays easy for users and organisations to generate and share content, individually or within communities. Social media sharing properties, such as Flickr, Facebook and PicasaWeb host billions of images and video, which have been annotated and shared among friends, or published in groups that cover a specific topic of interest. The fact that users annotate and comment on the content in the form of tags, ratings, preferences etc and that these are applied on a daily basis, gives this data source an extremely dynamic nature that reflects events and the evolution of community focus. Although current Web 2.0 applications allow and are based on annotations and feedback by the users, these are not sufficient for extracting this "hidden" knowledge, because they lack clear semantics and it is the combination of visual, textual and social context, which provides the ingredients for a more thorough understanding of social content. Therefore, there is a need for scalable and distributed approaches able to handle the mass amount of available data and generate an optimized 'Intelligence' layer, also called Collective Intelligence, that would enable the exploitation of the knowledge hidden in the user contributed content.

There already exist a number of approaches, which use user-contributed content in order to provide useful information for various applications. For example, mobile location information and uploaded content is used to monitor online traffic and generate traffic patterns in [1], connect citizens in Boston [2], share nature experience [3], discover travel patterns and provide travel advice [4][5], or communicate problems in a city [6]. The MIT Center for Collective Intelligence¹ is hosting a series of projects aiming at harnessing and using Collective Intelligence. These include the Climate Collaboratorium [7], which deals with climatic changes and the Collective Prediction effort, which tries to make accurate predictions about future events such as product sales, political events, and outcomes of medical treatments [8]. Collective Intelligence is also used in healthcare [9], while studies are conducted on its applications in today's organizations [10]. However, the main characteristic of such applications is that they are mostly based on collecting well-structured contributions through specific applications, on shallow statistical processing of the contributions and their visualization. Very few focus on analysis and on dealing with unstructured large-scale data, where an important source of knowledge is hidden.

Large-scale user contributions and more specifically, tags, which can be used to extract Collective Intelligence, suffer from a number of limitations, such as polysemy, lack of uniformity, and spam, thus not presenting directly an adequate solution to the problem of content organization. Therefore, how to manage, index and search for this content effectively and efficiently is becoming an increasingly important research topic. There have been many approaches dealing with the relevant tasks of tag refinement (to refine the unreliable user-provided tags) and automatic annotation or tagging, especially for user-contributed photos. In [11] the Content-based Annotation Refinement (CBAR) method re-ranks the tags of an image, reserving the top ones as the refined results. A more elaborated method [12] refines and enriches tags based on the visual and semantic consistency residing in social sites. Other algorithms perform automated tagging of an untagged image, either by building classifiers for individual semantic labels [13][14], or by learning relevance models between images and keywords [15][16]. The most frequently used method for automatic tagging is the semi-supervised graph-based level propagation technique [17], where a graph is constructed to model the relationship among individual images in terms of visual similarity. Such approaches, although they make use of the available content collections and their connections, they are trying to improve the quality of each independent user contributed content item and therefore they do not directly address the objective of extracting the hidden Collective Intelligence, which can be used in applications other than annotation and retrieval.

1 <http://cci.mit.edu/index.html>

Existing approaches towards extracting Collective Intelligence usually build upon restricted combinations from the available social media attributes. For example, in [18] geo-locations and tag information are used in order to generate representative city maps. In [19] tags and visual information together with geo-location are used for objects (e.g. monuments) and events extraction. Tags from Flickr images and timestamp information are used in [20] to form a chronologically ordered set of geographically referenced photos and distinguish locals from tourist travelling. The description of city cores can be derived automatically, by exploiting tag and location information [21]. The approach is able of distinguishing between administrative and vernacular uses of place names, thus avoiding the potential for confusion in the dispatch of emergency services.

But besides these combinations, user generated content can be viewed as a rich multi-modal source of information including attributes such as time, favorites and social connections. For example, beyond harnessing content and the surrounding tags or text, limited effort has been made to include the social patterns into the media analysis. The important aspect of fusion of modalities and different sources is currently lacking in existing Collective Intelligence applications. In this chapter novel techniques for exploiting these multiple layers of intelligence from user-contributed content are presented, which together constitute Collective Intelligence, a form of intelligence that emerges from the collaboration and competition among many individuals. The Collective Intelligence technologies to be described were developed in the context of the FP7 EU project WeKnowIt: Emerging, Collective Intelligence for personal, organisational and social use². User contributed content is analysed by integrating research and development in visual content analysis for localisation (Media Intelligence), tag clustering and Wikipedia ontology-based categorization (Mass Intelligence), analyzing social structures and user communities access rights (Social Intelligence) and event representation (Organisation Intelligence). The exploitation of the emerging Collective Intelligence results is showcased in two distinct case studies: an Emergency Response and a Consumers Social Group³ case study.

The chapter structure comprises an overall of nine sections. After this Introduction, the following section details the current state-of-the-art for each intelligence layer. The next four sections describe the developed technologies in each intelligence layer individually, namely the Media, Mass, Social and Organisational Intelligence layers. The seventh section presents the integration of the different technologies within one framework, namely the Integrated Collective Intelligence Framework (ICIF), which is exploited in the Emergency Response and Consumer Social Group use cases. The eighth section describes these two application scenarios where the presented techniques have been used together in order to leverage Collective Intelligence. The section also includes user evaluation results for the developed demonstrators. The last section contains conclusions and possibilities on building on top of the achieved Collective Intelligence results.

2. Background

The presented 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. Collective Intelligence methods can be classified based both on the number of the input modalities or layers that they employ in the analysis and the number of users contributing to the data. More specifically, the different Intelligence layers that contribute to Collective Intelligence can be classified to the following:

Media intelligence is the intelligence originated from digital content items (images, video, audio, text) and contextual information analysis, either provided by the user or pre-existing, and their merging. For this purpose, intelligent, automated content analysis techniques are used for different media to extract knowledge from the content itself. Since the amount of data is large and noisy, machine learning, data mining and information retrieval methods are used. Also the methods are able to fuse information from different sources/modalities, contextual information (e.g. time, location, and EXIF metadata), personal context (profile, preferences, etc.) and social context (tagging, ratings, group profiles, relevant content collections etc.).

Mass intelligence analyzes user feedback. Mass analysis enables input information clustering and ranking as well as information and event categorization. Also, bursts of information can be detected that may indicate potential events (emergency) and trend analysis and prediction. Facts and trends are recognized and modelled by interpreting user feedback on a large scale. For instance, a single road being blocked in a storm may not be very critical, but all access roads being blocked towards a hospital centre may be very critical in the case of an emergency.

² <http://www.weknowit.eu>

³ <http://weknowit.research.yahoo.com/csg/>

Social Intelligence is the exploitation of information about the social relations between members of a community. Nearly everything humans do, they do in a social context because they communicate, collaborate or in some other way interact with other people. Information about the various types of social relations may be represented in communication networks, friendship networks or organization charts.

Organisational Intelligence allows support of decision making through workflows exploiting the generated knowledge and taking into account existing procedures within an organisation. This is quite a departure from traditional methods where knowledge is produced by the individual knowledge worker and collected and integrated manually in knowledge based systems or organisational repositories.

2.1 Advances in Media Intelligence

Following current web socializing and multimedia changing user trends, most users nowadays upload, describe, geo-tag and localize their personal photos, based on previous personal or community content. Undoubtedly, the growth of such image collections and change of user behaviours created the need for intelligent algorithms, able to analyse mass heterogeneous multimedia content. In the following we focus on content-based retrieval systems whose aim is to identify the same object under various viewpoints within a large database and where in most cases local features extraction is no longer sufficient. Perhaps the most popular features used for object categorization are the SIFT features [22]. A typical example of using local patches around key-points for the retrieval of objects is [23]. A very well-known category of approaches is commonly known as “bag-of-words” model, which shares significant similarities with text retrieval approaches. This model has been adopted by Sivic and Zisserman [24]. Towards the need of lowering down this complexity, Chum et al. [25] propose a technique called Locality Sensitive Hashing and a random sketch of the set of visual words present in the image is used as an image representation. Another approach that uses a hashing scheme is presented in [26], where sets of feature vectors are indexed under their partial correspondences in sub-linear time. Checking the spatial consistency between the query image and the top retrieved images is adopted in the works [27] and [28], using the well-known RANSAC parameter estimation algorithm, introduced by Fischler and Bolles [29]. Geometric hashing has been introduced by Wolfson and Rigoutsos [30] for matching geometric features against a database.

2.2 Advances in Mass Intelligence

Collaborative Tagging is nowadays a standard feature of content sharing web applications enabling users to: (a) upload new, or bookmark existing content and, (b) annotate it by means of free-text keywords (tags). Created folksonomies in such Social Tagging Systems constitute a direct encoding of the views of a large number of users on how content items should be organized through a flexible annotation scheme (tagging). Detected tag clusters reveals relations between tags perceived by users as related to each other. To date, tag clustering has been dealt with either by conventional clustering algorithms, such as k-means [31] and Hierarchical Agglomerative Clustering [32], or by use of community detection methods [33]. Conventional clustering schemes are frequently troubled by two shortcomings: (a) the need for providing the number of clusters as input to the algorithm, and (b) their computational complexity. The new tag clustering scheme was designed with the above limitations in mind.

Created folksonomies and large-scale ontologies in addition to discovery of tag communities, can be used as significant resources in categorization. Accumulated knowledge helps to overcome the shortcoming of the classical approaches, like Support Vector Machines [34], Naïve Bayes [35], or Latent Semantic Analysis [36], as they all require training set of pre-classified documents. When a training set is not available, alternative approach should use available knowledge such as named entities, relationships between them and ontology schema to perform classification. In such approach named entities and relationship between them can be successfully used for term disambiguating and vocabulary unification (Bloehdorn and Hotho, 2004), or calculation of semantic relatedness [37]. Additionally, descriptions of neighbouring entities can enrich information about a classified document [38].

2.3 Advances in Social Intelligence

The usage of social intelligence proved to be useful in different areas. Winerman [39] shows that first information in crisis situation is often not provided by professionals but average citizens. She investigates how to exploit this knowledge to create official community-response grids.

Besides analysing the pure content of the provided information it can be helpful to use complementary data about sources, distribution and routing of information. Social network analysis provides tool sets. New ways of information acquisition require methods and tools to route, store and operate these data. Existing methods in access control (to govern these new ways of collaboration) do not really fit the new demands. Standards like XACML [40] cannot support self-organizing communities with distributed access rights management. Another example lacking self-organization is EPAL [41]. To provide intelligent information routing the identification of social groups and structures is a key requirement. As online social networks can be of large size, fast and scalable algorithms are necessary. Few methods are capable of identifying groups in networks with millions of nodes like the Label Propagation Algorithm by [42].

2.4 Advances in Organisational Intelligence

Events are understood as the occurrences in which humans participate. They may be very complex and a variety of aspects need to be considered. Models of events exist in various domains like the Eventory [43] system for journalism, the Event Ontology [44] as part of a music ontology framework, the ISO-standard of the International Committee for Documentation on a Conceptual Reference Model [45] for cultural heritage, the event markup language EventML [46] for news, the event calculus [47] for knowledge representation, the Semantic-syntactic Video Model [48] and Video Event Representation Language (VERL) [49] for video data, and the event model E [50] for event-based multimedia applications. From this related work we have derived that event aspects such as time and space, objects and persons involved, as well as mereological, causal, and correlative relationships between events, and interpretations of events have to be considered. The Event-Model-F [51] we have developed provides full support for these requirements. It advances the current state of the art by its full support for causality, correlation, and interpretation of events. The ER Log merging and management (WERL) application makes use of the Event-Model-F and is related to the general domain of C4I software, which stands for Command and Control Systems and Components. Among the numerous solutions available in the market that target at a wide range of applications such as military operations and surveillance, WERL is most closely related to ER incident management solutions. Three well-known solutions are: the Atlas incident management system (Aims) [52], the Emergency Command System [53], and the Bristol City Council map-based application [54]. The main focus of such software solutions is the support for information sharing and communication, as well as task and asset management. However, they lack the reusability and shareability of information that WERL offers thanks to its Event-Model-F-based representation, as well as the semantic enrichment capabilities of WERL.

3. Media Intelligence

In spite of the multitude of current intelligent, automated visual content analysis activities, there is still a lack of appropriate outlets for presenting high quality research in the prolific and prerequisite field of multimedia content retrieval. Thus, the goal of developing intelligent, automated content analysis and retrieval techniques for different media to extract knowledge from the content itself remains a major research task; still image retrieval turns out to be one of the most exciting and fastest growing research areas in the field of multimedia technology that aids towards the taming of those needs.

In principle, designing Collective Intelligence with a user-centred approach requires the involvement of users from the very beginning, as it is fundamental to understand the reality of what are people doing, how, when, and why. A user-centred design approach often works by trying to answer typical questions like who are the users, which are the user tasks and goals, what information do the users need and so on. Therefore, the undoubtedly growth of multimedia content collections and the change of user behaviours have created the need for fast, robust and efficient Media Intelligence algorithms, able to analyse large-scale diverse and heterogeneous visual content. More specifically, the popularity of social networks and web-based personal image collections has resulted to a continuously growing volume of publicly available photos and videos. Users are uploading, describing, tagging and annotating their photos, whereas recently they also geo-tag the location they were taken. In addition, a heavily increasing percentage of people are using the Internet to find and provide additional information with respect to past or forthcoming events or actions.

In this framework, the main scope of Media Intelligence tackles the area of such visual content interpretation. Media Intelligence suggests a content-based retrieval approach and focus is being given on a visual retrieval and localisation framework of content applicable to large web collections that may become extremely useful in pre-, during or post-travelling user activities. Motivated by this

observation, a fast and robust retrieval system is being proposed [80], based on the popular bag-of-words model. More specifically, Speeded-Up Robust Features (SURF) have been selected to capture the visual properties of digital images and a visual vocabulary is created, through which images are efficiently represented. Geometric constraints on the image features are then taken into account, facilitating more accurate retrieval in comparison to traditional approaches. The performance of the proposed method is evaluated through the development of a web-based image retrieval application, that yields a geographic position estimation about a query image, exploiting already geo-tagged datasets. The latter aids users to identify so-called Points of Interest (POIs) and famous landmarks (i.e. in a “what to see?” concept), as well as additional background info about them together with interesting activities and events (i.e. in a “what to do?” concept).

In the following, we present the VIRaL tool⁴, a web-based tool used to identify and localise similar multimedia content under different viewpoints, applicable to any functionality that involves a still image search and retrieval task. The main research principle of VIRaL exploits the fact that typical metadata usually contain a free text description together with some representative user-generated tags. In some cases, some metadata are related to the geographical position of the image taken (a.k.a. geo-tags). A geo-tag consists of the actual geographic coordinates, i.e. the longitude and the latitude, and is either extracted automatically through GPS or is manually defined by the user.

Initially, in order to represent the visual content of any given digital still image, a set of interest points is selected and visual features are extracted locally from their surrounding area. Since the goal is to choose scale invariant interest points, their localisation is carried out on a Gaussian scale-space. Using SURF features has been proven to achieve high repeatability and distinctiveness, whereas their extraction speed is very fast, when compared e.g. with SIFT features [21]. An example of the extracted SURF features is depicted in Fig. 3.1.




Fig. 3.1. SURF features extraction.

To further understand the notion of a visual vocabulary, one should consider it as an equivalent to a typical language vocabulary, with an image corresponding to a part of a text. In the same way that text may be decomposed to a set of words, an image can also be decomposed to a set of visual words. Then, in order to compare two images, their corresponding visual words may be compared. Thus, it is interesting to create a visual vocabulary in such a way that parts of images could be meaningfully assigned to visual words. Fig. 3.2 depicts two regions of interest extracted from two different images,

which correspond to the same visual word. The visual vocabulary is presented as the Voronoi cells of the clustered space of visual words. We should note here that due to their polysemy, visual words cannot be as accurate as natural language words.



Fig. 3.2. Regions of interest extracted from two different images that correspond to the same visual word.

To create the visual vocabulary, a clustering process is followed. More specifically, the well-known K-means clustering algorithm [55] is applied on the SURF descriptors corresponding to a very large number of points of interest. If the number of the points to be clustered is significantly large, clustering using the K-means algorithm becomes a very slow task. For example, clustering of 5M of points (which are typically extracted from 10K of images) requires a few days of processing. However, to efficiently deal with large scale retrieval problems, the size of the vocabulary should be in the order of a few tenths of thousands of visual words [27][56]. Thus, in order to rapidly create an appropriate vocabulary, the clustering process is performed on a smaller subset, carefully selected to contain the most representative images. Finally, after constructing the visual vocabulary, each image has to be represented with a description that captures its relation to all the words of it.

The process of querying an image database without and with a visual vocabulary is depicted in Fig. 3.3. In the first case the comparison of the local descriptors is performed immediately for two images and after exhaustive comparisons in the whole database, the closest regions are found. In the latter case, for every image of the database all points have been assigned to appropriate visual words of the visual vocabulary. Thus, for a new query, its points have to be assigned to the closest visual words of the vocabulary. After this process, two images are considered to be similar if their points are assigned to similar visual words.

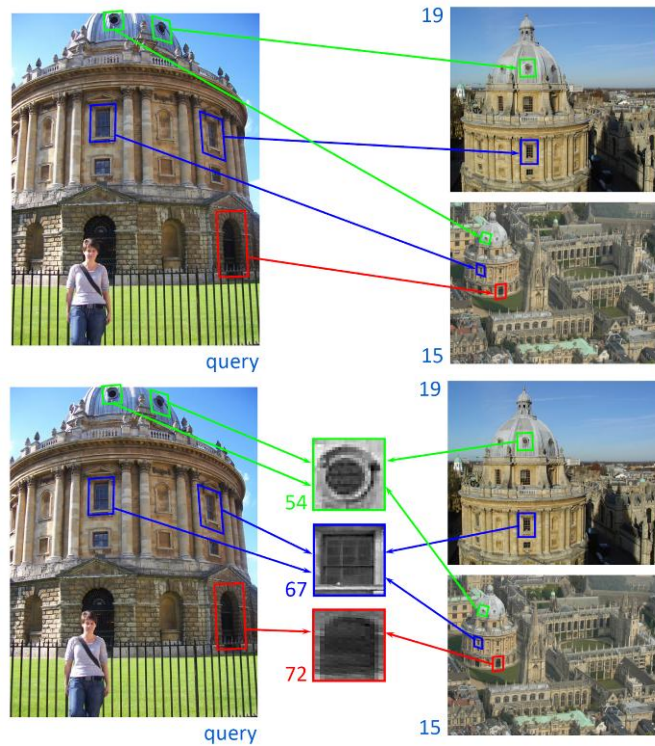


Fig. 3.3. Querying an image database without and with a visual vocabulary.

When a user query reaches the system, then the local low-level features are extracted from the query image and the model vector is computed. Then the similarity of the query model vector with all database model vectors is computed using an inverted-file structure to speed up the process, and the N most similar images, the images with the highest such value, are either returned to the user as similar, or become candidates for geometric consistency check using the RANSAC algorithm [57], [80]. Finally, regarding geo-tag estimation, if a user issues a query containing a landmark image against a large database of geo-tagged images, then most probably the top-retrieved results will contain the actual landmark that the query image depicts. Those correctly retrieved images are expected to have near identical geo-tag values. Little variance is expected since geo-data can be defined by users and also because the same building may be photographed by different distances (using appropriate camera lenses). However, the estimated geo-tag for the initial query image is expected to be within the larger consistent subset of the result images (Fig. 3.4, 3.5 and 3.6).

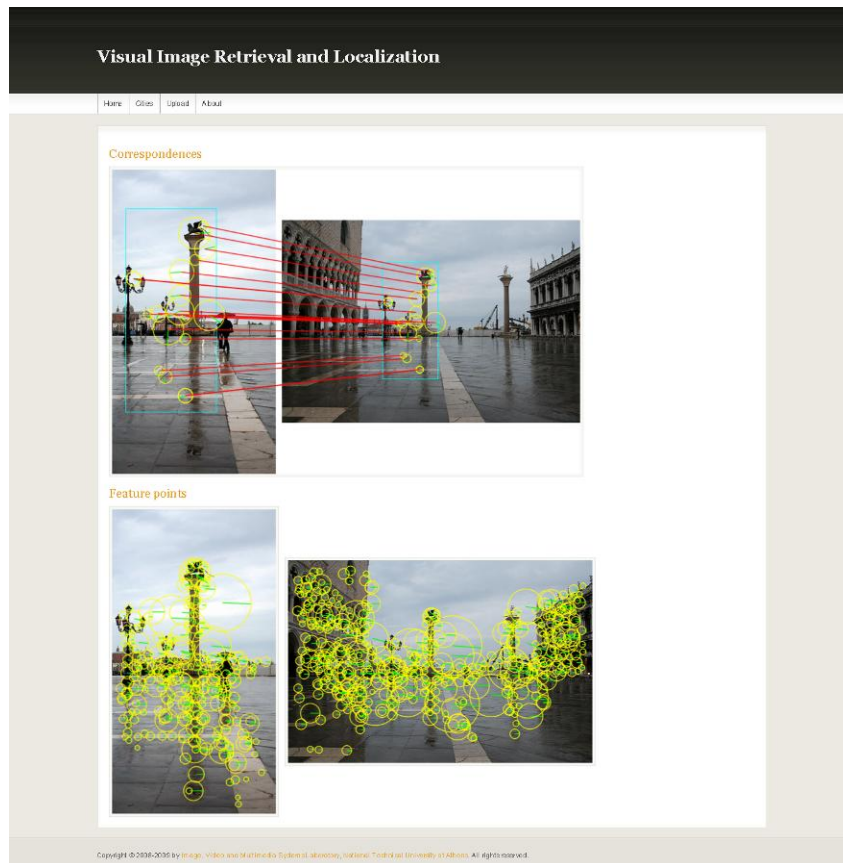


Fig. 3.4. Correspondence of images.

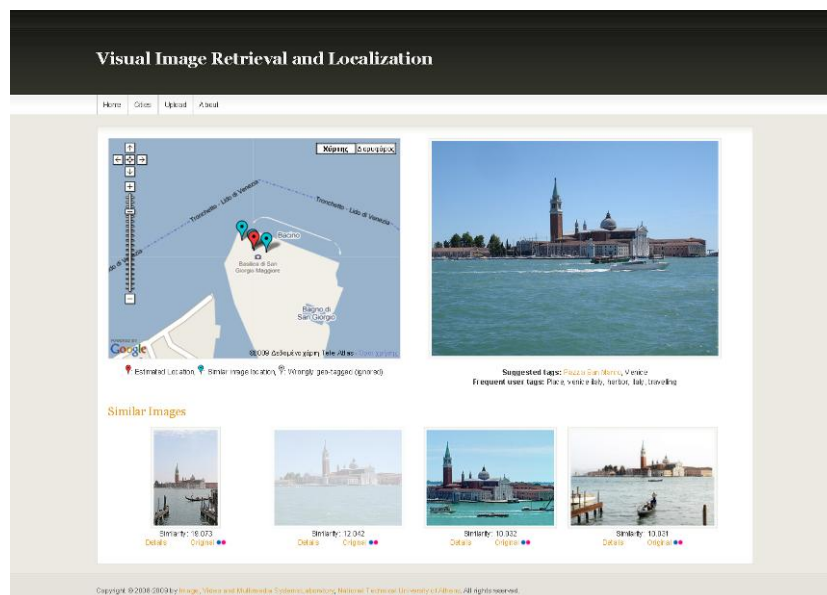


Fig. 3.5. Geo-tag estimation.

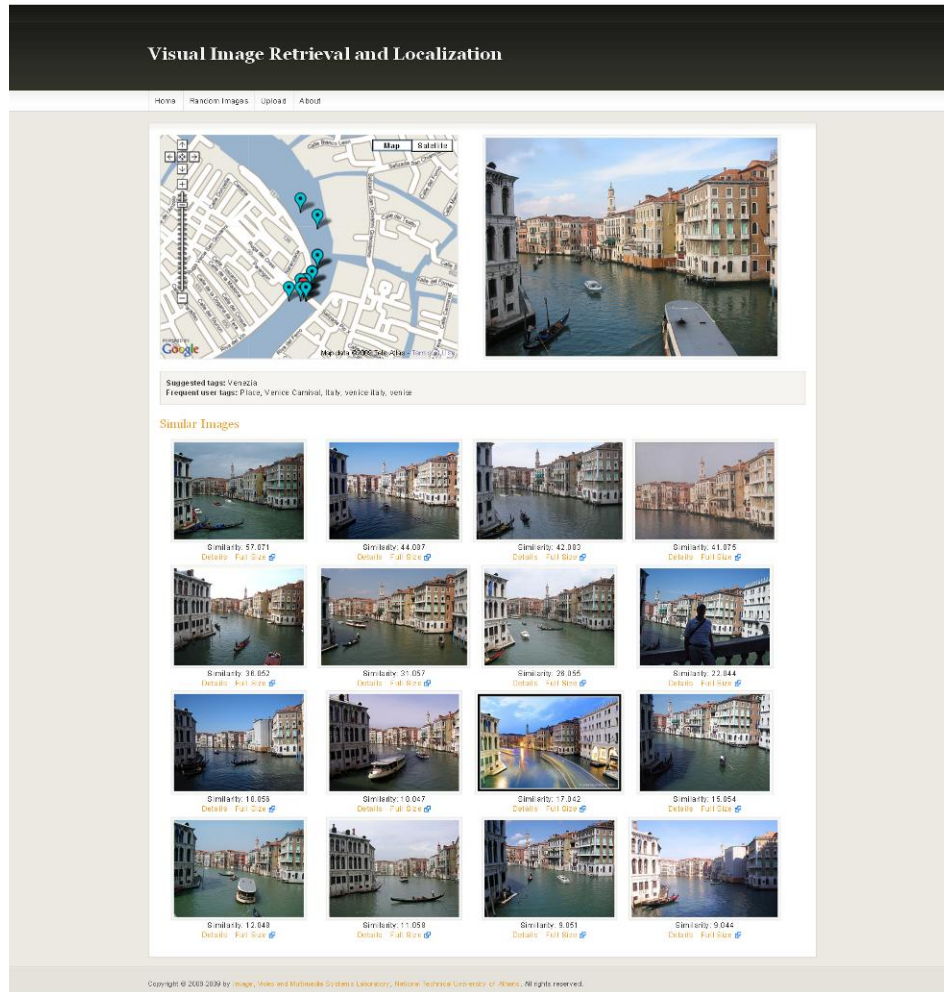


Fig. 3.6. Consistent subset of result images.

The above described methodology has been evaluated on a challenging 1M urban image dataset, namely European Cities 1M⁵. It consists of a total of 1.037.574 geo-tagged images from 22 European cities, which is a subset of the dataset used in the VIRaL tool. In order to acquire detailed information on the evaluation process, the reader is encouraged to consult [80] or [81].

4. Mass Intelligence

The vast amount of available and produced information in the current web requires new approaches to categorization and clustering in order to help users in efficient navigation and information finding. With the analysis of tag networks and training-less ontology-based categorization, large-scale user generated content became the centre of interest for Mass Intelligence. Nevertheless, Mass Intelligence does not only concentrate on such information. It can consume both massive direct user input as well as aggregated inputs already processed by other intelligence layers. Processing of inputs or joining results of the other intelligence layers creates basis for the Collective Intelligence. The example approaches presented below also cooperate with other intelligences to create new values and solutions, like hybrid image clustering that joins tag community detection with clustering based on visual features created in Media Intelligence.

4.1 Community detection on tag graphs

Folksonomies comprise three types of entities, namely users, resources and tags, as well as the associations among them [58]. Tag clustering involves a process that groups the tags in a way such that members of the same tag cluster are perceived by users as related to each other. Despite the subjectivity of users involved in judging the degree of relatedness between tags, tag clusters are expected to correspond to meaningful topic areas, which can be useful in a series of tasks [59], such as

information exploration and navigation, automatic content annotation, user profiling, content clustering and tag recommendation.

The proposed scheme builds upon the notion of (μ, ε) -cores introduced in [60]. The original algorithm, referred to as SCAN, suffers from two problems. First, it needs two parameters, namely μ and ε , to be provided as input. Second, it leaves a substantial number of nodes unassigned to clusters. As a result, its utility is limited in IR tasks such as tag recommendation. For that reason, our scheme conducts an efficient iterative search over the parameter space (μ, ε) in order to discover cores for multiple values of the parameters. Finally, the identified cores are expanded by maximizing a local measure of modularity [61] in order to increase the number of nodes that are assigned to communities and to allow for overlap among communities. The scheme is described in detail in [62] and its three steps are briefly explained below:

Core set discovery. The definition of (μ, ε) -cores is based on the concepts of structural similarity, ε -neighborhood and direct structure reachability. The structural similarity between two nodes is defined as:

$$\sigma(u, w) = \frac{|\Gamma(u) \cap \Gamma(w)|}{\sqrt{|\Gamma(u)| \cdot |\Gamma(w)|}},$$

where $\Gamma(u)$ is the structure of node u : $\Gamma(u) = \{w \in V \mid (u, w) \in E\} \cup \{u\}$. Then, the ε -neighborhood of a node is the subset of its structure containing only the nodes that are at least ε -similar with the node, i.e. have similarity equal to or higher than ε . A node is called a (μ, ε) -core if its ε -neighborhood contains at least μ nodes and a node belonging to the ε -neighborhood of such a core is said to be directly structure reachable from it. Eventually, community seed sets are extracted by identifying the (μ, ε) -cores of a network and attaching to each one of them the nodes that are structure reachable with them.

Parameter space exploration. One issue that is not addressed in [60] pertains to the selection of parameters μ and ε . Setting a high value for ε (the maximum possible value is 1) will render the core detection step very eclectic, i.e. few (μ, ε) -cores will be detected. Moreover, higher values for μ will also result in the detection of fewer cores (for instance, all nodes with degree lower than μ will be excluded from the core selection process). For that reason, we employ an iterative scheme, in which the community seed set selection operation is carried out multiple times with different values of μ and ε so that a meaningful subspace of these two parameters is thoroughly explored and the respective (μ, ε) -cores are detected. The scan of the parameter space starts from high μ and ε values, moves logarithmically towards lower μ values, then lowers ε by a small step and starts again from the a high μ . It terminates as soon as it reaches the lowest meaningful (μ, ε) values.

Core set expansion. Once the community seed sets have been identified by the above process, a core set expansion step is carried out in order to enrich existing communities with more relevant nodes. It achieves this by starting a local exploration process with the goal of maximizing a local quality measure, namely subgraph modularity [61]:

$$M(S) = \frac{ind(S)}{outd(S)},$$

where $ind(S)$ stands for the number of within-subgraph connections for subgraph S , and $outd(S)$ stands for the number of connections from subgraph nodes to the rest of the graph.

Example of the tag community around tag ‘computers’ is presented in Fig. 4.1.

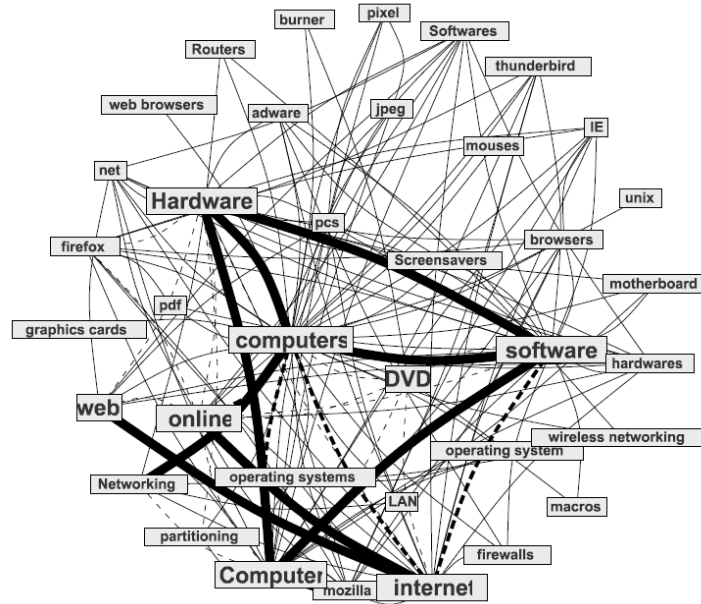


Fig. 4.1. Tag community around tag “computers”.

In qualitative evaluation our method produced much more meaningful tag clusters compared to CNM [63], which contained few gigantic clusters and many small ones, and comparable clusters, but richer in terms of tag coverage than the ones produced by SCAN [60].

Qualitative evaluation was based on subjective assessment of the derived tag communities and an implicit evaluation by using the derived clusters for tag recommendation and measuring the achieved performance on historical tagging data from three different tagging sources (Delicious, Flickr, and BibSonomy). Complete results are reported in detail in [62].

The process of extracting tag clusters from massive user tagging leads to promising results, however it is solely based on the statics of tag usage. As a result, the extracted tag clusters may often contain irrelevant tags (low precision) or miss related tags (low recall). Tag cluster precision can be improved by exploiting large-scale semantic resources, such as WordNet and Wikipedia, in the way presented in the following section. Cluster recall can be improved by propagating tag descriptions of photos to other photos that are visually very similar to them. Visual similarity can be established with the analysis tools provided by the Media Intelligence layer. In that way, a synergy between two intelligence layers is established in order to improve the quality of the analysis result.

This Collective Intelligence driven idea resulted in a Hybrid Image clustering approach joining Media and Mass Intelligences. An image similarity graph was built encoding both the visual and the tag similarity between images of the collection.

Experiments and evaluation of the hybrid image clustering are described in details in [82]. According to the performed user study results, used image clustering approaches are characterized by very high precision scores ($\geq 90\%$). Visual-only clusterings are characterized by superior precision ($\approx 98\%$), but suffer from low recall. Tag-only clusterings behave in an IR-complementary way, yielding higher recall rates at lower precision. This allows creating of an image cluster that captures pictures of the same landmarks taken from very different spots. Evaluation was conducted on a set of 128,714 geotagged images located within the metropolitan area of Barcelona. Sample image cluster detected by the hybrid approach is presented in Fig. 4.2.

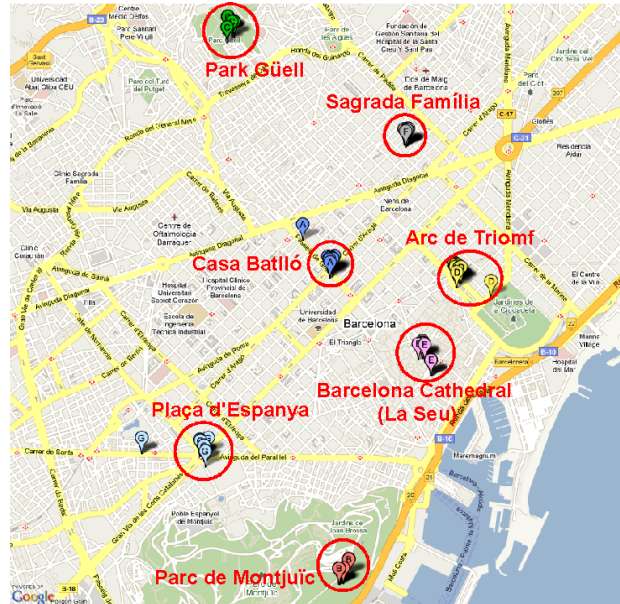


Fig. 4.2. Example Barcelona landmarks that were identified as image clusters by hybrid approach.

The result of this Collective Intelligence instance can be further exploited by the Organisational Intelligence layer, namely the WeKnowIt Emergency Response Log merger and manager, for improved log entry indexing and retrieval.

4.2 Ontology-based classification

The method relies on the domain knowledge represented in the form of an ontology to perform the categorization task. It concentrates on the recognized named entities and relationships in the document text to measure the semantic similarity of the created thematic graph to the categories, defined as ontology fragments, and perform the categorization. In the proposed text categorization method the ontology effectively becomes the classifier. As a result, it overcomes the limitation of classical approach, and does not require a training set of pre-classified documents. Instead, it allows defining categories as fragments of ontological knowledge.

Our categorization algorithm consists of three main steps: (1) construction of the semantic graph, (2) selection and analysis of the thematic graph, and (3) categorization of the selected thematic graph. The complete approach is presented in [64] and [65]. The approach uses ontology created from Wikipedia due to richness of represented domains, high number of interconnected entities, and included categorization scheme.

Semantic graph construction requires identification of named entities (based on entity phrases and labels known to ontology), relationship extraction with shallow NLP and connectivity inducement that utilizes ontology as background knowledge. Analysis of a sample sentence is presented in Fig. 4.2.

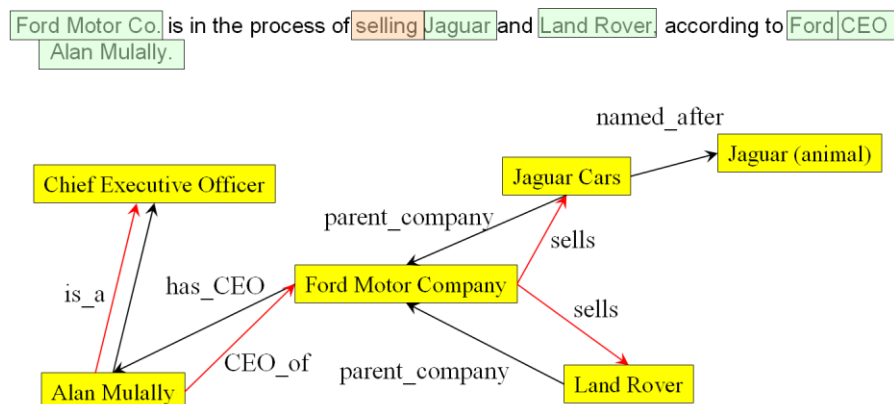


Fig. 4.3. Sentence with marked phrases and created semantic graph.

Thematic graph is selected after identification of created components in semantic graph and finding most authoritative entities using HITS algorithm [66]. The dominant thematic graph is further taken for categorization. It is based on discovered entities and initial categories assigned in ontology.

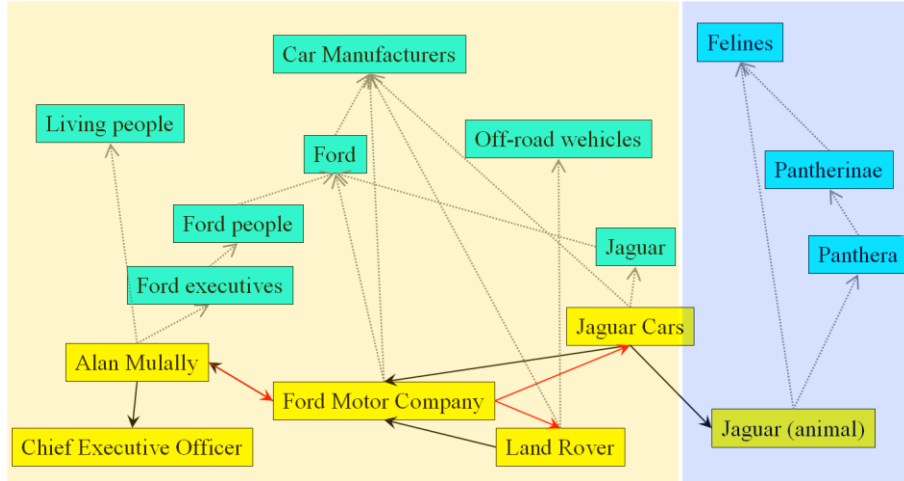


Fig. 4.4. Entities with category trees from ontology.

Separate category trees for identified entities also can perform disambiguation function. As presented in Fig. 4.3 entities matched from the sample sentence belong to two category trees: automotive (core) and animals (side category). Concentrating on core entities and dominating category tree, establishes categorization of the whole document.

The document including presented sample sentence was categorized by the presented method as Wikipedia categories: “Hybrid vehicle” and “Automobile”. The semantic graph created from the whole news document is presented in Fig. 4.4. Most authoritative entities are highlighted.

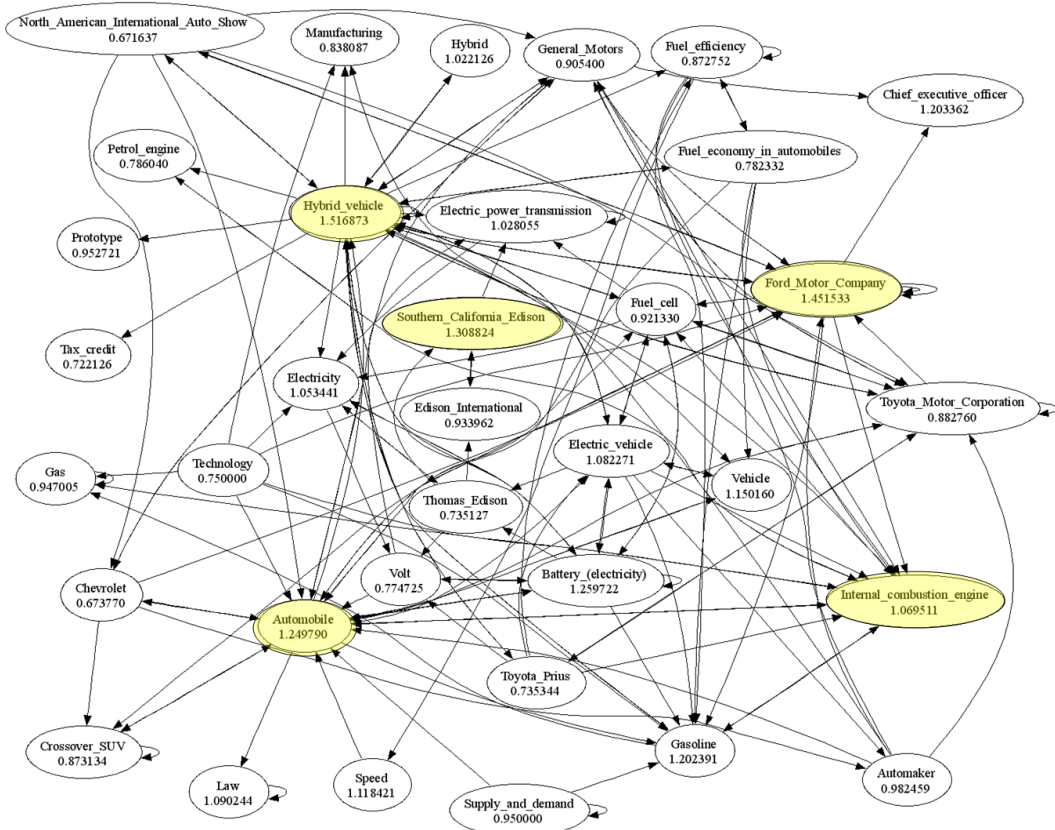


Fig. 4.5. Semantic graph from news document about hybrid vehicle announced by Ford.

In the performed experiments on news from CNN (www.cnn.com) RSS feeds and subset of the Reuters RCV1 corpora [67] we compared our method with Naïve Bayes from BOW implementation

[68] and SVM from WEKA package [69]. The proposed method reached over 85% accuracy without the need of training set.

5. Social Intelligence

The analysis of the structures of social networks reveals information on general properties of the relations, on the hierarchical composition of a network and on the roles and positions of people. One example for the general analysis of network properties we conducted is the analysis of the communication on Question & Answer platforms. This analysis revealed the topic-specific answer behaviour. The knowledge about the different structures of communication explained the varying performance of expert identification algorithms. Expert identification is used to guide users to domain experts for one or several topics or to provide users information on the trustworthiness of fellow users. Gaining an understanding of how people interact or relate to other persons is the first step in extracting knowledge from social networks.

For example the knowledge of 'natural groups' of similar people is a key issue for many real-world applications. If the members of a person's social group are known, new social applications become feasible. One of those applications is the social emergency alert service (EAS) which activates members of the social group of a person in need in a case of emergency [83][84]. This service promises to facilitate faster help than the status quo. The social emergency alert system consists of software clients for smart phones and a server component. By analysing the users' communication (voice and text) social groups can be identified and used for broadcasting help request. A feasibility assessment showed that there is a good chance to receive help from a near-by friend [83].

Social networks can be very large. The networks of popular online communities like MySpace and Facebook have hundreds of millions of users. Therefore, scalability is one of the major problems for network analysis techniques. A community identification algorithm that is able to process huge networks is the randomized greedy modularity clustering algorithm (RG) we proposed in [70]. Modularity is a measure for the quality of a graph clustering (i.e. the decomposition of the set of nodes into non-overlapping groups) introduced by Newman and Girvan [71]. RG is an algorithm that tries to identify groups by maximizing the modularity of a graph partition. This algorithm is explicitly designed for natural networks like social networks that have many structural similar substructures. The agglomerative hierarchical algorithm exploits these redundancies by using a randomization strategy that finds with a high probability but little effort nodes that belong together to iteratively build a near optimal dendrogram. Processing a network with about 5 million nodes and 40 million edges takes less than 10 seconds on standard desktop computers. Another algorithm we developed is optimized for the processing of huge dynamic datasets [72]. The algorithm is based on restricted random walks that are incrementally updated. The incremental update process avoids that for every change of the data set all data needs to be reprocessed.

The EAS is a good example for Collective Intelligence: It makes use of network data from the different layers of intelligence to reason about social groups.

5.1 Protecting Virtual Communities

Analysing online communities to support their members is one element of social intelligence. Another task is, to create a safe environment where people can interact. For almost any IT system, access control is an important issue. Especially for modern, state-of-the-art social networking platforms like Yahoo!⁶, Facebook⁷ and Flickr!⁸, a flexible, reliable, manageable and easy way of access control is important. Personal data and media as well as social interactions have to be protected. Users want to define in detail, who can access their media, e.g. holiday photos. They also want to differentiate access with regard to the social groups they belong in a convenient way.

The same principle can be applied to two use case scenarios, Emergency Response Scenario and a Consumer Group Scenario. Both scenarios rely on media objects like images, videos, text documents and audio streams. It is obvious, that an access rights concept is necessary to allow or disallow access to these resources.

In traditional approaches, access rights are mainly modelled by 3-tuples of the form (user, permission, object). Access is granted, if the current system state matches one of such 3-tuples in the access rights

6 <http://www.yahoo.com>, last accessed 30.6.2010

7 <http://www.facebook.com>, last accessed 30.6.2010

8 <http://www.flickr.com>, last accessed 30.6.2010

facts (“facts”). Giving an example, the 3-tuple fact (alice, read, pic1.jpg) allows a user with an account named Alice to access a file pic1.jpg with read access.

More generally, the 3-tuple (U, P, O) is an element in the space $U \times P \times O$. U hereby is the set of users or subjects, sometimes also called principals. Subjects are entities, typically users or programs, executing in a system on behalf of a real world user [73]. P is the set of permissions. Typical elements of this set are read, write, execute and delete. The semantic of an access right is usually given to it by the developers of the application or software system. O represents the set of objects the access right applies to. Typical objects can be files and devices. It is important to notice, that user accounts can be subjects as well as objects, depending whether they act active (do something) or passive (something is done with them).

Two conclusions are obvious: By this design, a very fine-grained facts modelling is possible. Secondly, it is obvious, that it is practically impossible to manage a more complex system with several thousand objects and lots of users by defining these triples individually. The access rights facts base would be huge, and any change has a quite high probability to affect several access rights which all must be changed manually. This leads to a high error rate by access rights not correctly set.

To avoid this, two major approaches have been introduced. First, several elements are grouped in sets. For example, all user accounts of a department are grouped assigned to one group. The same can be applied to groups of access rights and objects. This grouping then allows defining facts on them and no longer on individual elements. Many approaches using this technique can be found [74][75].

Another improvement are the concept of authoritative roles and access control lists (ACL). Roles combine permissions and objects (e.g. “read file A”, “write all files in dir B”). Roles become an intermediate layer, a container holding objects and permission combinations. Roles are assigned then to users, allowing them to perform operations described by their assigned roles. Usually the role approach is combined with grouping. Roles allow to abstract access rights from individuals by modelling rights in abstract roles. If a role of a user changes his new access rights can be modified easily by changing his roles.

ACLs do work similar as roles, but use a different perspective: ACLs combine users and access rights to access control lists, which are then assigned to objects. We see immediately, that this is the same approach as roles, applied this time on objects and not users.

The common RBAC model [76] (Role-based Access Model), for example, is an example of the first extension. In RBAC access rights - the two latter elements of the triple - are collected in so called roles. In our example a very simplistic role could be defined by role1, allowing to read Pic1.jpg. Roles are then assigned to users. Obviously, the RBAC and the triple model have the same expressive power. RBAC collects permissions and objects in roles and assigns them to users, while the basic triple model directly models triples.

It is important to notice, that these approaches do not extend the functional expressive power of an access rights model. The same result could still be retrieved by the triples model. The extensions focus on usability and manageability. They relay all on the same *authorization model*, thus the (u, p, o) triple.

5.2 The Community Design Language

We suggest a different approach allowing defining not only policies but defining and modifying the underlying authorization model (Fig. 5.1).




Fig. 5.1. Layers of an Authorization System.

The Community Design Language (CDL) is a formal language to define (1) the model, (2) the policies and, (3) deduce access rights.

The model describes the possible conditions which can be used to formulate policies in. Policies are the high-level definition of the conditions which grant or deny access. The access rights are the deduced combinations of entities which allow or deny access.

For a better understanding, we provide an example: In the model it is defined, which circumstances and conditions may be used to formulate policies. To model a “traditional” *u-p-o* model, the conditions user, permission and object must be defined. If additionally the access time shall be introduced as possible policy condition, it is defined in the model. Let us define for this example, users, permissions, objects and access time in the model. A policy defines – based on the model – if access shall be granted. In our example, all users of the group scientist may read all files through the intranet. This is a policy. The deduced access rights is then the fact, that Alice (as member of the group scientists) can read the file pic1.jpg from the network intranet. This example is presented in the Community Design Language in Table 5.1.

This approach enables a convenient and powerful approach to define access models, policies and access rights.

The CDL is designed to make use of external knowledge sources to decide upon access. This allows the integration of other intelligence layers. For example, a CDL policy can be defined to grant access to images identified been taken in a specific geographical region.

Table 5.1. This table shows the example presented in the text in the formal language CDL.

<p>Model definition:</p> <p>CREATE SETS users, permissions, files, networks; CREATE SETS usergroups; CREATE RELATION usingroups (users, usergroups);</p> <p>Define policies:</p> <p>CREATE ACCESSCONDITION ac: ([users].usergroups IN “scientists”, [permissions] IN “read”, [networks] IN “intranet”);</p> <p>Fill facts in database:</p> <p>CREATE ELEMENTS users: { Alice }, permissions: { read }, files: { pic1.jpg }, networks: { intranet }, usergroups: { scientists };</p> <p>CREATE LINK usingroups: {(Alice, scientists)};</p> <p>Perform access check:</p> <p>CHECK ACCESS (users=Alice, permissions=read, files=pic1.jpg, network=intranet); // result is true</p>

6. Organisational Intelligence

Organisational Intelligence is the ability of an organisation to understand and to leverage knowledge that is relevant to its goals and purpose. As a consequence, it is the goal of Organisational Intelligence to bring the right piece of knowledge, at the right time to the right person, in order to support decision making to best accomplish the organisations purpose. For instance, in emergency response, typically several professional entities are involved such as emergency hotline, police department, fire department, and emergency control center. All these entities need to exchange event descriptions like the one above. However, they typically use different systems and applications with their own proprietary data models for events. Using the formal Event-Model-F [51] instead, these systems can commonly represent and effectively communicate event descriptions.

The Event-Model-F bases on the foundational ontology DOLCE+DnS UltraLight (DUL)⁹ and provides a set of ontology design patterns to represent the different relations of events as derived from the related work in Section 2. The participation of objects in events is implemented by the participation pattern. This pattern also provides for modelling the absolute time and location of events and objects. The mereology pattern, causality pattern, and correlation pattern implement the structural relationships between events. The mereology pattern allows representing composition of events along temporal,

9 <http://wiki.loa-cnr.it/index.php/LoaWiki:DOLCE-UltraLite>

spatial, and temporal-spatial relations of events and objects. The documentation pattern provides for annotating events, e.g., by media or sensory data. It can be seamlessly linked with other ontologies like the Multimedia Metadata Ontology [77] for precisely describing digital media data. Finally, the interpretation pattern supports different event interpretations.

6.1 Event Log Merger Application

The Event-Model-F is used in the ER Log merging and management (WERL) application. WERL addresses the problems arising in reviewing and searching through the logs that are produced by different members of the ER personnel. The disparate log entries are automatically merged and represented on the basis of Event-Model-F. Furthermore, semantic information is extracted from their text in order to enable a concise view of the ER log content.

Commonly, a log file contains information pertaining to the documented incident and the log creator, as well as a set of timestamped log entries, each of which documents a message communicated between some members of the ER personnel and an associated action. Thus, two granularities of events are defined, the high-level emergency incident (e.g. a fire incident in a factory) and the specific actions taken by the ER personnel, which are considered as sub-events of the high-level incident (composition pattern). Depending on the granularity of the event, different documentation properties are attached to it (documentation pattern). Furthermore, log entries are described based on their location and temporal attributes as well as the involvement of specific individuals to them (participation pattern).

In order to derive several of the aforementioned log attributes (location, person names, etc.), the recorded log text is undergoing a semantic enrichment process. As training data is not available to learn the extraction patterns and inconsistencies are observed in the log entries, in terms of linguistic and syntactic style, the extraction processes does not rely on natural language patterns, but applies a knowledge-intensive approach. This requires that quality resources (gazetteers/taxonomies) are available, containing the desired named-entities likely to be found in the logs. For instance, for emergency incidents there is a requirement to identify fine-grained locations (i.e. at the street level). In addition to the extraction of location and person names, prominent key phrases and ER-specific terms are extracted from the text.

The semantically enriched log entries are surfaced to the professional users through the WERL front-end. The front-end of WERL provides online filtering capabilities for facilitating the interactive exploration of the available log entries. A snapshot of the application main screen is provided in Fig. 6.1. At the top, a slider-based time filter is available that enables the examination of a particular time interval of the incident. In addition, standard full text search capabilities are provided for retrieving only the subset of log entries that are relevant to the input query. Most importantly, there is a series of four semantic filters that summarise the main entities found in the log files by the text annotation component of the system. Thus, it is possible to view only the log entries that are related to a particular location, person name, or significant keyword or ER acronym. The presentation of all identified semantic entities in these lists can provide the ER user with a quick overview of the semantic content of the log file.

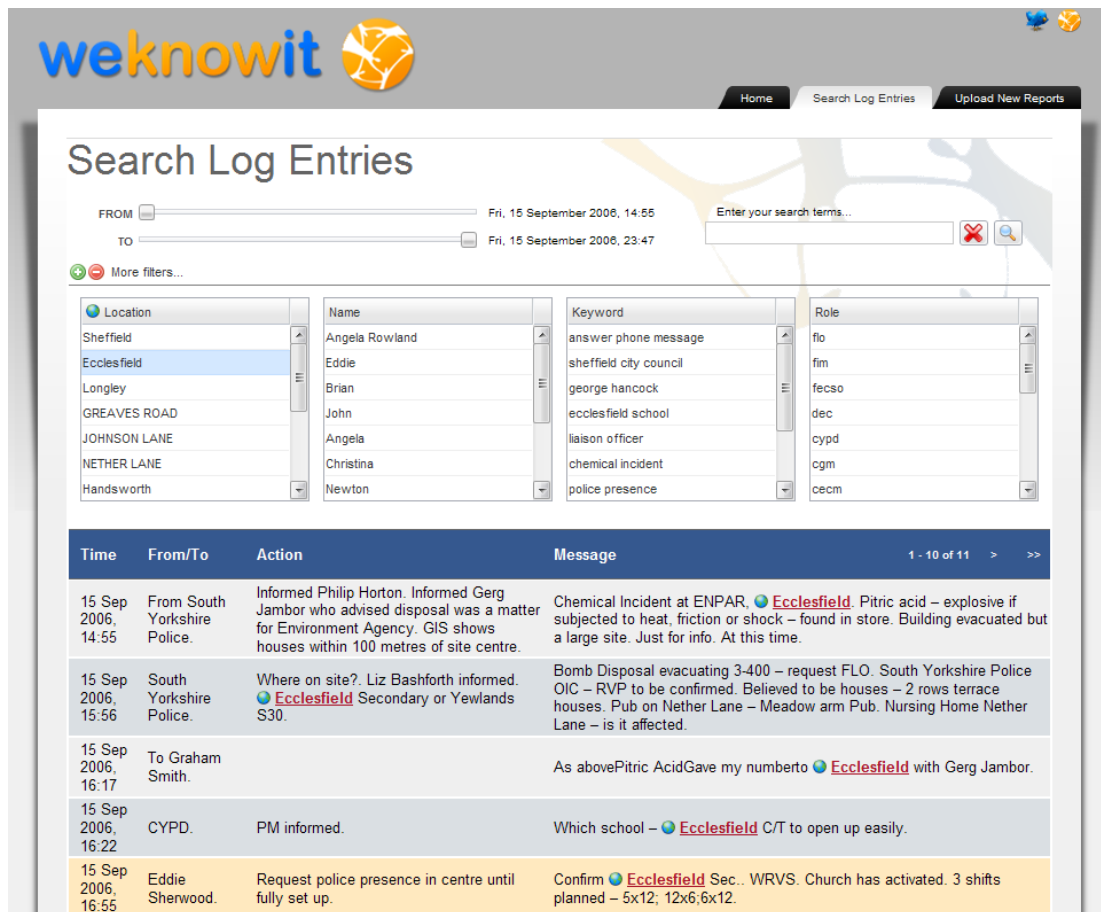


Fig. 6.1. Snapshot of the WERL front-end main screen.

Another significant feature of WERL is the presentation of provenance information and the possibility to filter based on the provenance of log entries. Beside each log entry there is a marker indicating its origin. At the bottom of the log entry list, there is an associated legend, which can also be used for filtering based on the log entry provenance. In that way, it is possible to inspect only the log entries produced by a particular log creator, thus gaining insight into his/her perspective of the incident. In larger scale incidents, involving many members of ER personnel coming from different organisations (e.g. fire department, police), it is expected that more sophisticated provenance mechanisms will be necessary, e.g. provenance by organisation, unit, role in the organisation, etc. Finally, WERL provides a map-based view of the log entries based on the automatic identification of fine-grained location information from their text.

6.2 Sharing Event Descriptions with SemaPlorer

Besides the use of the Event-Model-F in the log merger application it is also used in the SemaPlorer++ application, an extension of the SemaPlorer [78] application, for creating and sharing event descriptions. A domain specific ontology on emergency incidents provided by the Sheffield City Council (SCC) has been developed and is used within the SemaPlorer++ application. The SemaPlorer application allows its users to interact with events on the map view as shown in Fig. 6.2. An ontology browser is located on the left hand side with an emergency response ontology loaded. The emergency events defined in the SCC ontology are listed in a tree-structure of the ontology browser. It enables the user to easily create event descriptions by clicking on a concept in the ontology, i.e., clicking on the emergency event IAEP_Major_Industrial_Fire in the ontology browser representing a major industrial fire, and dragging and dropping it on the map.

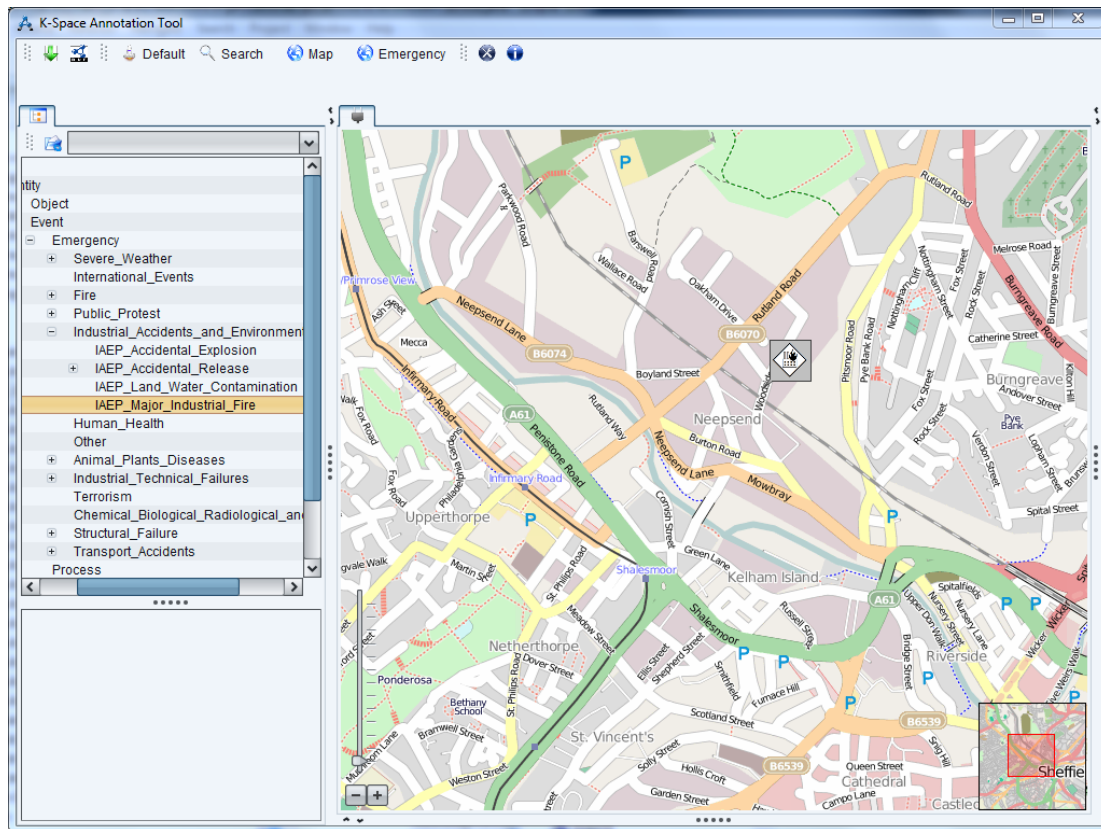


Fig. 6.2. Screenshot of the SemaPlorer++ Application for Creating and Sharing Emergency Response Event Descriptions.

Once the user has placed an event description on the map, an instance of the event participation pattern of the Event-Model-F with the information about the event is created. It comprises the time and location of the event and a default object participating in the event. In addition, an instance of a specialization of the event documentation pattern is created and connected with the event participation pattern. This specialization of the documentation pattern allows storing additional information about the event such as a title, written location name, description, and documenting pictures. The example in Fig. 6.2 shows an icon on the city map of Sheffield. It represents a major industry fire that happened in a bakery in Sheffield.

6.3. Application of the Event-Model-F to Tourism and Sports

With the Event-Model-F, we can create and exchange sophisticated descriptions of real world events. It has been formally specified in Description Logics [85] using the Web Ontology Language [86]. We have verified its validity by using the build-in reasoning tools of the ontology engineering tool Protégé [87]. Besides the use of the Event-Model-F in the two applications above for emergency response, it can also be applied to arbitrary other domains that need to represent events as occurrences in which humans participate. Different examples for applying the Event-Model-F include soccer games and tourism and are available at: <http://west.uni-koblenz.de/eventmodel>.

The soccer example models an entire soccer game, i.e., the first half-time and second half-time. Different events happen during the game such as a foul and goal. The soccer example makes full use of all patterns of the Event-Model-F, namely participation, causality, correlation, composition, and interpretation. It further shows how a domain specific ontology can be embedded and used to describe the events happening during a soccer game. A tourism example models a two-day weekend trip. Three people are participating in this trip. On the first day, there is a sub-event dinner. On the second day are two sub-events, a visit to a museum and a sight. The tourism example applies different patterns of the Event-Model-F such as participation and composition.

7. Integrated Collective Intelligence Framework

The technologies described in the previous sections can be used alone. But its cooperation brings real value to the end users. To this end a new Collective Intelligence methodological approach 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.

From a technical perspective, the role of Integrated Collective Intelligence Framework (ICIF) is to achieve the synergy effect ("Collective Intelligence") by combining the results of work from the different services into a cohesive system.

In order to achieve this, the following technical tasks have been fulfilled:

- specification and implementation of the overall software and hardware architecture,
- preparation of a set of commonly used objects and implementation of a storage component capable of storing them,
- integration of services (software components) based on the common model and API.

7.1 Collective Intelligence methodology

As can be seen in the intelligence layers chapters, the presented 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 the Wikipedia-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 proposed Collective Intelligence approach moves a step further; instead of a mere concatenation of the different layers intelligent methods, it imposes a pairwise 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. Geo-tagging through visual and tag analysis yields better localisation results, while WERL can achieve improved log entry indexing and retrieval when incorporating the Collective Intelligence tag clustering method.

Fig. 7.1 depicts the proposed Collective Intelligence approach, which 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

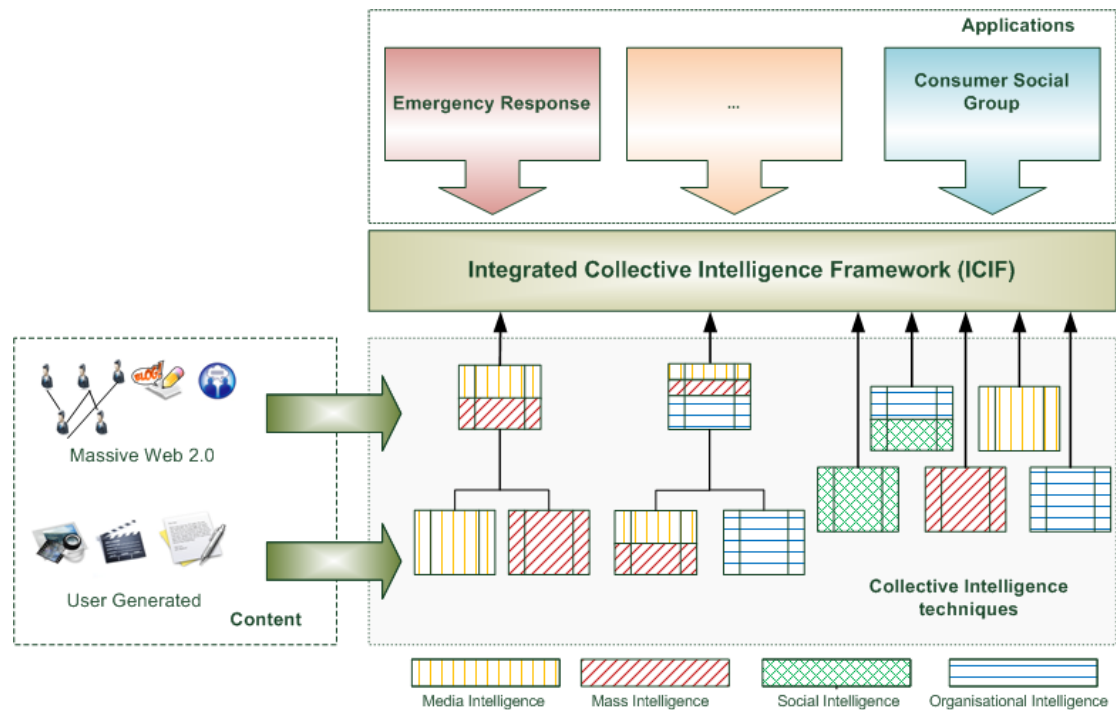


Fig. 7.1. Collective Intelligence Methodology

7.2 Architecture and Integration

The architecture of the ICIF is based on the idea of loosely-coupled components. Cooperation of the components is realized via the registry of the OSGi¹⁰ framework. The connections between the components can be specified programmatically or declaratively (using enterprise integration patterns). Some functionality provided by the ICIF are exposed via REST API allowing external applications to use ICIF's services.

Taking into account the massive calculations performed by some components (e.g. visual analysis), a scalable architecture has been prepared. An Enterprise Service Bus (ESB) component is responsible for passing messages between services. Additionally, some "heavyweight" components can run on different machines and be accessed remotely.

All components of the system are provided as OSGi bundles and integrated within the ICIF architecture.

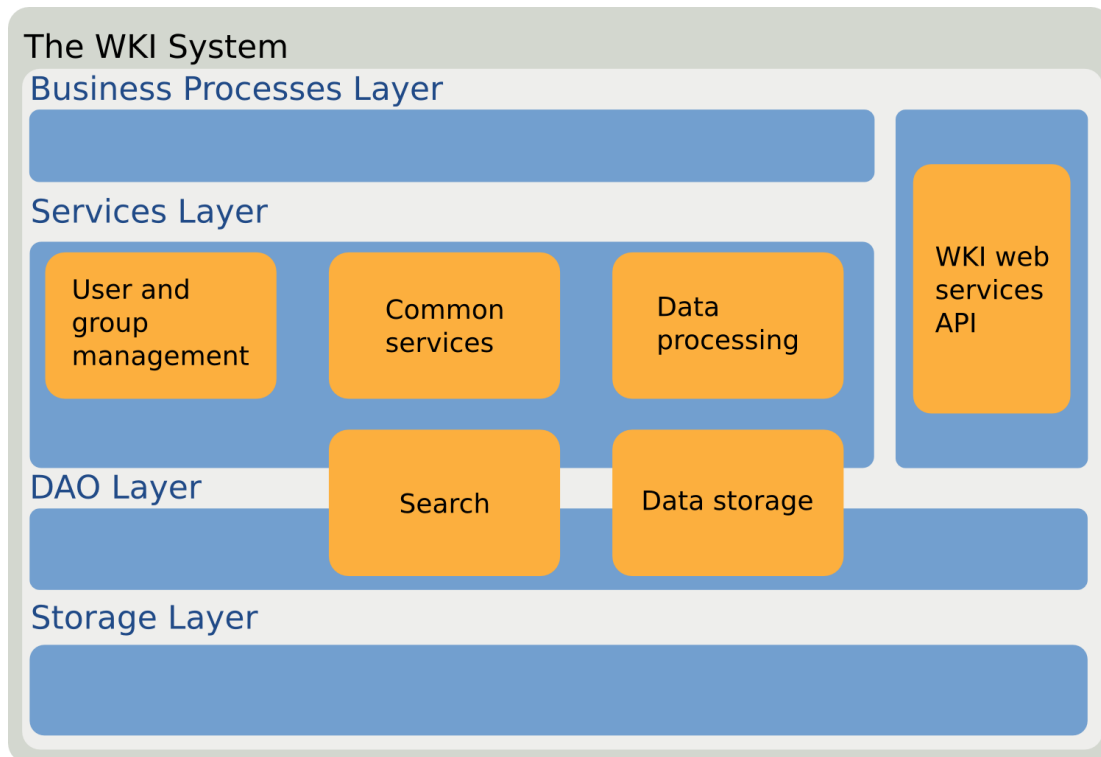


Fig. 7.2. Layers and groups of modules within the WKI System.

Fig. 7.2 presents how the services, described in previous sections are grouped into (conceptual) modules and deployed within the ICIF platform.

The ICIF is written in Java and uses many open-source enterprise ready technologies including Fuse ESB (integration platform), Apache Camel (integration framework), Apache Karaf (OSGi framework), Apache CXF (web services framework), Spring (Java/J2EE application platform), Lucene (fulltext indexing and search engine) and other¹¹. All of them are mature, and proved their stability in enterprise projects. No vendor-specific products are used, and open communication standards are utilized. Currently, more than 15 services (of varying granularity), including the ones presented in this chapter and others developed in the framework of the WeKnowIt project, are integrated within the ICIF. Table 7.1 provides exemplary services belonging to each of the aforementioned conceptual groups. Note that each group can contain services from different Intelligence Layers. For example group “Data processing” includes services from Media and Mass Intelligence.

Table 7.1. Services of the WeKnowIt System.

Group	Services examples
User and group management	WP4_CommunityDesignLanguage, WP5_GroupManagement
Common services	WP6_CommonsModel
Data processing	WP2_TextClassification, WP2_TextClustering, WP2_VisualAnalysis, WP2_SpeechIndexing, WP3_SpamDetectionService, WP3_LocalTagCommunityDetectionService
Search	WP6_DataStorage, WP2_SearchInSpeech
Data storage	WP6_DataStorage

Some of the services are accessible on their own, while other are combined in order to perform multiple operations in response to one call of the client applications. For example, a simple act of file upload triggers a chain of actions within the ICIF platform. First, a mime-type of a file is determined. Provided that the uploaded file is an image the following services are executed:

- WP4CommunityDesignLanguage service determines user rights to upload a file and stores permissions of the newly uploaded file after it was successfully processed,

¹¹ <http://fusesource.com/products/enterprise-service-mix/>, <http://camel.apache.org/>, <http://karaf.apache.org/>, <http://cxf.apache.org/>, <http://www.springframework.org/>, <http://lucene.apache.org/>

- WP2VisualAnalysis service determines the geo-location of a picture and returns tags that are likely associated with it,
- WP2TagNormalization service matches the tags to a domain ontology concepts and adds some annotations to metadata,
- WP6DataStorage stores the file with all generated metadata.

As a result, a file, along with extracted and generated metadata, is stored within the system thanks to which various applications can benefit from it.

7.2 Hybrid Storage

The storage of the ICIF (called WKI DS) is capable of storing both files and business objects (entities that are exchanged between the services of the system). It can be divided into three types of storages:

- file storage (Hadoop DFS),
- triple store (Jena), and
- object storage (NeoDatis).

All these storages are realized using open-source technologies. The pluggable architecture allows the actual implementation of each storage to be changed. For example, it is possible to replace Jena with some other triple store, if such need arises.

The idea behind the redundant storages for business objects (triple store and object database) is the following. On the one hand, the applications built on top of the ICIF require flexible access to data. This requirement is satisfied by the SPARQL endpoint of the triple store. On the other hand, a typical web application most frequently performs a number of operations during a certain period of time for which the triple stores are yet not designed. In order to make many operations (i.e. simple CRUD operations) perform much faster, the second type of storage - object database - is used.

There are two obvious problems with this approach though. First, there is redundancy of data that is duplicated among both storages. At the same time, there is necessity of mapping between objects and triples.

Regarding data redundancy, it is handled internally by the WKI DS component, which stores some additional information that allows for matching entities from object storage with graphs from triple store (this information is stored in separate relational database that is used only internally). The API of the WKI DS does not allow for operations on single triples in order to keep content of both storages synchronized. Insertions, updates and removals of objects are taken care of by the internal mechanism, which guarantees the consistency of data in both storages.

The mapping of business objects to triples is performed by an internal mechanism that is implemented using Jenabean library. The objects that the mapping mechanism can transform into triples (and the other way round) belong to a Common Model, which is a set of objects created based, on the common ontology that is used throughout the framework.

The WKI DS storage can be fine-tuned to meet the needs of an application. The object storage can be used as a full copy of triple store, but can be also configured to store only a part of its data (thus serving as a cache).

Fig. 7.3 presents this scenario. It is worth to notice that the mapping from triples to objects is completely transparent to the client ("Event service" in this example). In fact, the client is even unaware of the dichotomous structure of the WKI DS and operates solely on Java objects.

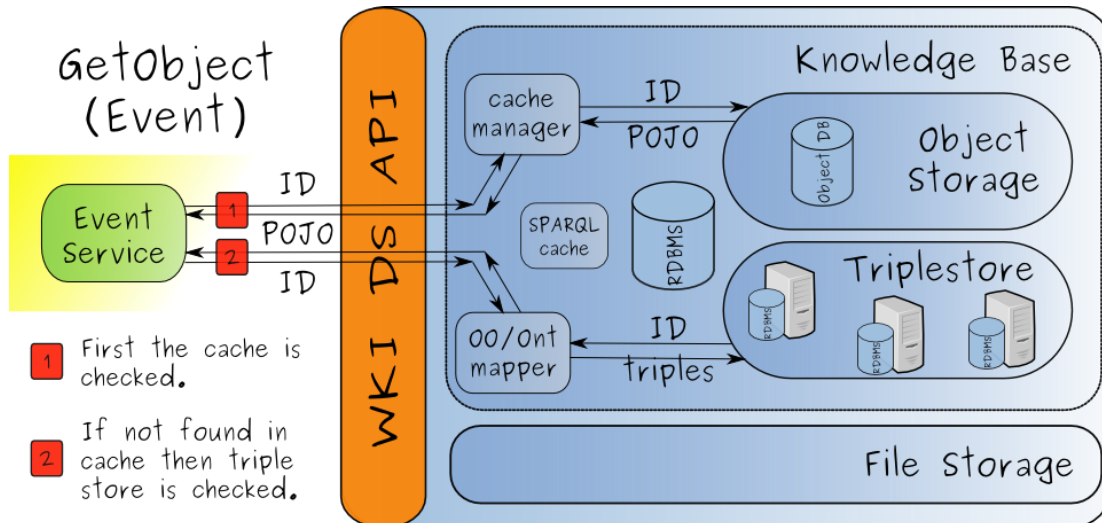


Fig. 7.3. Retrieval of business object from the WKI DS hybrid storage.

The existence of two storages allows for further improvements of the application's performance apart from benefiting from fast access to business objects guaranteed by the object database. An additional effort is being put into making the system even more efficient. Transformation of some of the SPARQL queries directed at the triple store into native queries of the object storage may result in a significant performance boost. The first proof of concept (implemented using a fixed set of queries from the Berlin SPARQL benchmark¹²) is very promising, but the whole idea still needs a lot of polishing.

7.3. Development Environment

Integration of work of independent development teams require good communication and coordination of efforts. It can be improved by the usage of proper open-source tools¹³:

- Source Code Repository (SVN),
- Continuous Integration server (Hudson) with code quality checking tools (Sonar),
- Artifacts Repository (Nexus),
- Bug Tracker (Mantis),
- WebDav server.

¹² <http://www4.wiwi.fu-berlin.de/bizer/BerlinSPARQLBenchmark/>

¹³ <http://www.mantisbt.org/>, <http://nexus.sonatype.org/>, <https://hudson.dev.java.net/>, <http://subversion.tigris.org/>, <http://www.sonarsource.org/>, <http://webdav.org/>

Fig. 7.4 presents a typical flow of development activities and role of different tools involved - from SVN checkout, through bug report, up to patch commit.

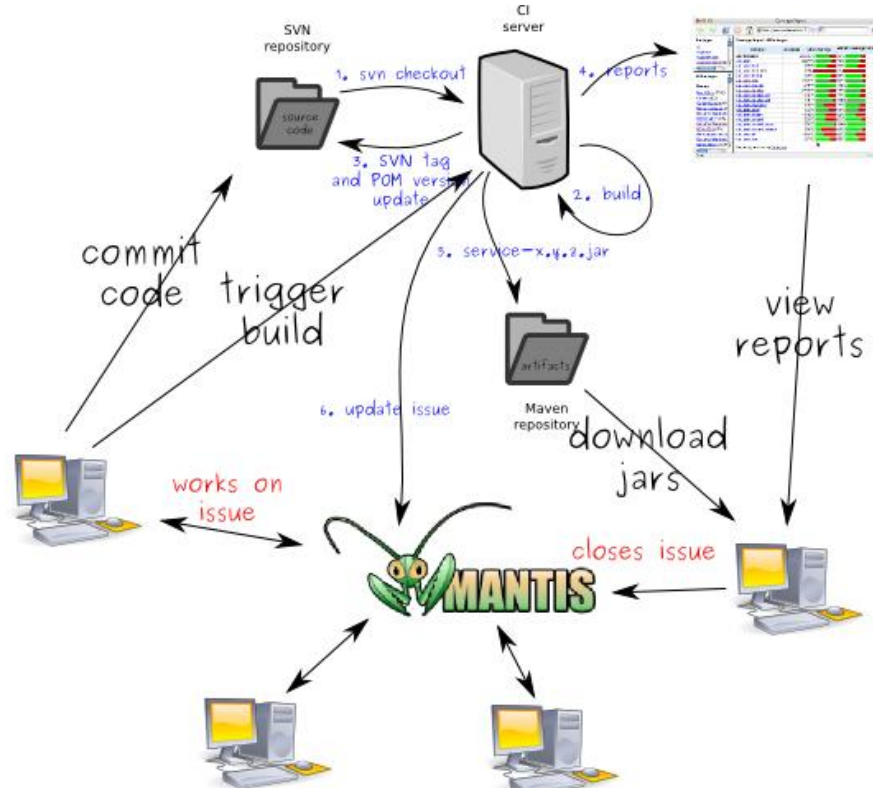


Fig. 7.4. Development environment of the Integrated Collective Intelligence Framework.

7.4 Summary

The ICIF platform brings into life Collective Intelligence by providing a runtime environment that can be used to combine services and achieve synergy effect. The first prototypes of an Emergency Response and a Consumer Social Group scenario prove this approach feasible. The hybrid storage can be used with different applications that require both fast business objects access and execution of complex SPARQL queries. Its unique capabilities open new perspectives for business applications that are powered by triple stores.

8. Applications

The techniques described in the previous sections have been integrated in two different scenarios in order to depict the feasibility of harnessing and producing Collective Intelligence. In an Emergency Response (ER) scenario (Fig. 8.1), upon an emergency event (e.g. fire, flood, etc.) a user logs in to the application and is able of capturing the event and contextualize it with metadata, e.g. tags. The Media Intelligence VIRaL localisation method is then used to automatically add location information to the image, by analysing massive input from Flickr. Thus, even if GPS is not activated or available, the uploaded image can be enriched with geo-localisation information. In the Emergency Responders side, Mass Intelligence techniques are applied in the user contributions, which are also enriched by relevant content from Web 2.0 sources. The community detection technique is applied on the received tag volume, so as to induce clusters of the received image data. Moreover, an ER domain ontology is used for the classification of the available tags. Meanwhile, if a user of the application is in danger, she is able of activating her connections in her social group, by utilizing the Social Intelligence Emergency Alert service. Finally, after the event, in the headquarters of the Emergency Planning team, the personnel is able of monitoring and reviewing the team's reports and actions by the use of the Organisational Intelligence ER log merger. In overall, the different analysis layers that are used in the different stages of the event contribute altogether to the leveraging of Collective Intelligence, which in turn yields better handling of such Emergency Response events.

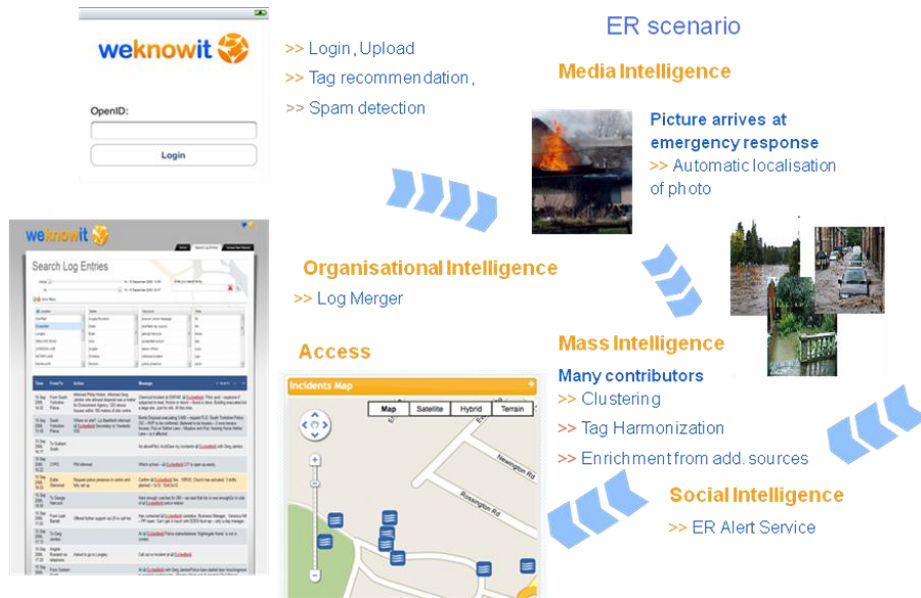


Fig. 8.1. Emergency Response scenario.

In a different scenario, the presented techniques can analyse user contributed content from various sources to help travellers discover essential information about what to see and do on travel or one-day cultural trip events. More specifically, by making use of automatically extracted Collective Intelligence results, it assists users in a travel exploration experience by identifying points of interest (POIs), ranking and prioritizing them (according to the most popular places and user's profiles) and by presenting them along with additional background information aggregated from different sources. The scenario contains three main parts: a Travel Preparation, a Mobile Guidance, and a Post travel stage (Fig. 8.2). During the travel preparation part, users need a tool, which is able to provide them with information about the different candidate places to be explored and visited.

The Travel Preparation tool, in the form of a web application, provides the relevant information that users need to prepare their travel, e.g. information about the locations, multimedia content, and points of interest. This information is aggregated with content coming from different sources (e.g. Wikipedia), is ranked according to trends to add value to user's experience, and is presented to the user by the use of Mass Intelligence clustering techniques. Fig. 8.3 depicts a snapshot from the online available travel preparation tool¹⁴.

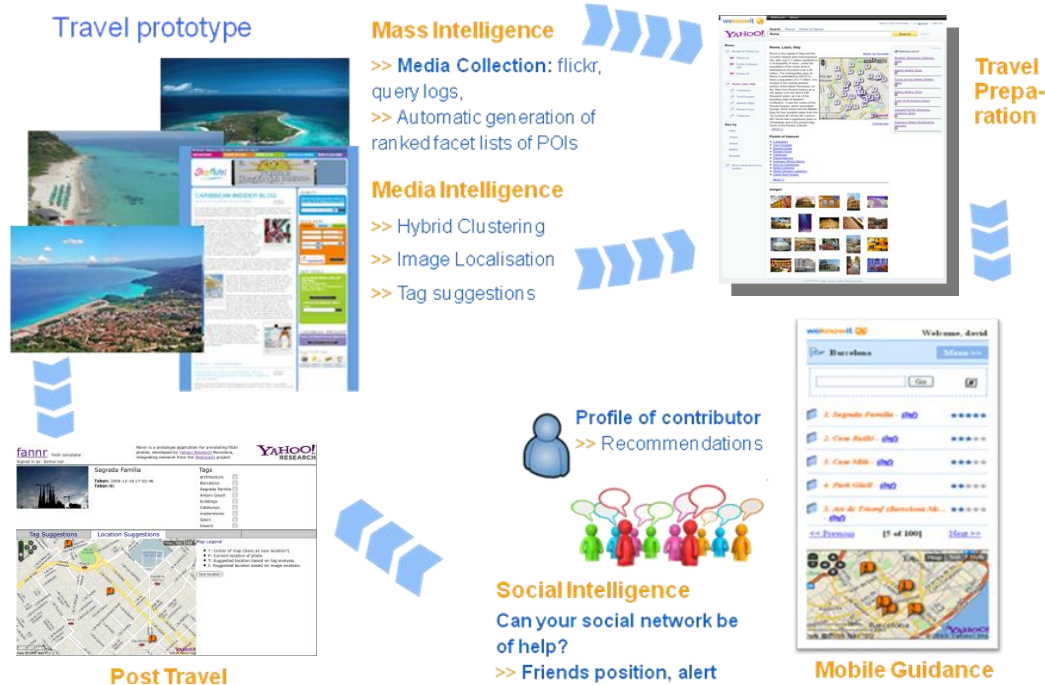


Fig. 8.2. Consumer Social Group scenario.

In the second part, Mobile Guidance, users perform the planned trip. With the features offered by mobile devices and the developed application, users are able to access relevant information about their physical environment and search for new events or points of interest. They are also able of using Social Intelligence techniques to find the position of their social connections, or even notify them in case of emergency. Users can also take pictures and record videos of the places they are visiting, which are likely to correspond to the ones chosen in the travel preparation phase. These data, generated by users, such as images and videos, comments, ratings and notes, can then be added into the system where all this content is stored and shared with other users, enriching the repository of information for other users of the system.

Finally, in a post travel application, the user can exploit the Media Intelligence image localisation technique, as well as the Mass Intelligence tag suggestion methods to enrich metadata in her data collection and also automatically geo-locate her photos.

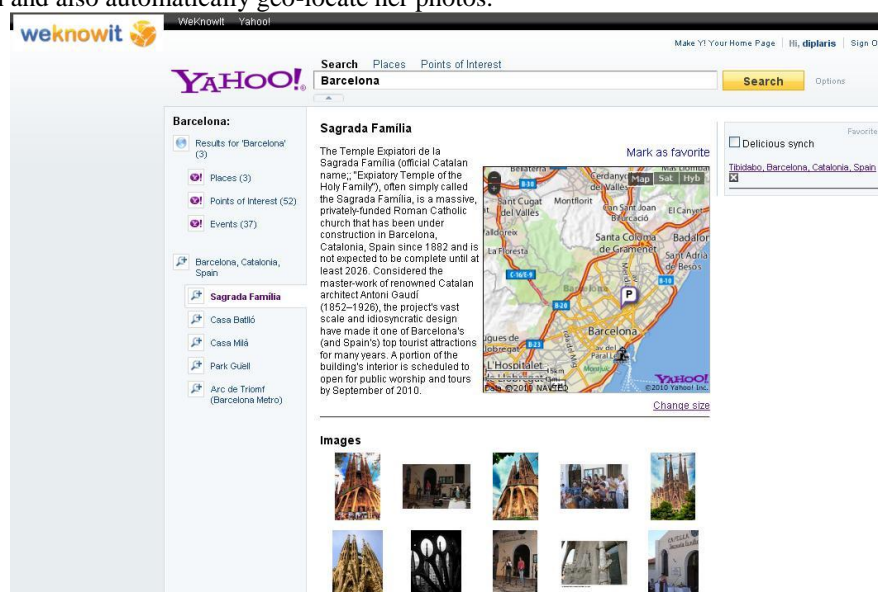


Fig. 8.3. Travel preparation tool snapshot.

8.1 Evaluation

The Emergency Response and Travel scenarios have been developed in the form of demonstrators. As the two demonstrators are harnessing Collective Intelligence in a different manner, different evaluation approaches and metrics are exploited for each case accordingly.

For the ER demonstrator two interfaces were evaluated during the evaluation runs – the access interface, used by citizens and ER experts to explore incidents and the upload interface, used to provide information to the system. The demonstrator was evaluated by two groups. Six Emergency Response experts evaluated the intelligent access interface to the Emergency Response demonstrator and a smaller group of experts evaluated the intelligent upload interface to the demonstrator. Additionally, 12 citizens evaluated both the upload and access interfaces. All evaluations were run in the city of Sheffield.

The desktop and mobile components of the travel demonstrator were evaluated separately. Two evaluation runs were carried out for each one of the demonstration modules, comprising 15 and 21 end users for the desktop module, while 4 and 22 end users evaluated the mobile guidance module accordingly.

The evaluations were split between lab exercises and field trials. Users were given tasks typical of the usage of the respective demonstrators. Following exposure to the demonstrators, the user response to the prototypes was evaluated through questionnaires or discussed in consensus meetings. These materials were analysed in order to assess the added value of Collective Intelligence and to identify recommendations for the future enhancement of the demonstrators. The evaluations addressed several dimensions: usability, complexity, efficiency, responsiveness, satisfaction etc. Feedback was also collected in a non-structured format in order to get explicit recommendations on how the demonstrators can be improved.

8.1.1 Evaluation of Emergency Response demonstrator

The evaluation focused on both whether target users were able to make use of the application to carry out the essential tasks and whether their performance in these tasks was better than the current status quo. To assess the functionality of the ER demonstrator, users were asked to carry out the tasks they would typically be involved in and assess the interface on the basis of these tasks. Thus, citizens were asked to upload an image and to use the intelligent access interface to determine information related to an incident they had previously witnessed in the city. The time taken, title, description and tags applied to the images were noted and the response to the interface was gathered. The ER experts were asked to primarily evaluate the access interface, although a secondary evaluation assessed the functionality of the upload interface.

The task for the ER experts was to gather information about and distinguish between three events which were co-occurring in the city. In both the access cases a further goal of the evaluation was to assess how the Collective Intelligence implemented by WeKnowIt impacted on the ability of citizens and ER experts to make sense of the information provided to them. Therefore, evaluations were carried out with three different interfaces in order to measure this impact, namely a) information gathered from incoming calls from the general public, b) raw data in the form of a simple web interface which allowed the users to see the images and comments in a serial manner, and c) the ER demonstrator. A number of images and comments were collected corresponding to the severity of the incident. ER experts were then asked to use the interfaces to build an understanding of each of the incidents, their location and their corresponding location. The citizen participants had access to the same data and interfaces but were instead asked to find out more information about one of the incidents (one with medium severity) on the basis of a small amount of information.

To measure the performance (as opposed to the efficiency) of the participants, a simple scoring system was used. The answers given by the participants were compared to the information used to generate the evaluation data on a one-to-one basis and given a score of 1 if the participant gave the same answer and 0 otherwise. The scoring was made for each incident, for the location, severity and type of incident. Thus each answer was scored out of 9 (9 meaning that the answer was completely correct and 0 meaning that the answer was completely incorrect). The scoring was carried out by a single person without reference to the condition to prevent biasing. Fig. 8.4 shows the mean scores per condition for the access evaluations.




Fig. 8.4. ER - Mean Score by Condition for Experts.

As Fig. 8.4 shows, the mean scores were generally high. As expected the score arising from the comments condition alone were lowest and the scores achieved in the ER demonstrator condition were marginally higher than those for the images (7.8 for ER demonstrator and 7.7 for Images). Given that the information present in the two interfaces is the same, this result is not surprising. The scores for the comments are lower; the near equality of the scores for the images and ER demonstrator conditions is reflective of the power that the image has for this user group.

Fig. 8.5 shows the average score achieved with each interface for the citizens.




Fig. 8.5: ER – Average score by condition for Citizens.

The difference in performance between the conditions was less pronounced in the citizens case than that for the ER experts, though this is largely due to the simpler task that the citizen was asked to perform. Again, however, it can be seen that the ER demonstrator receives the highest score overall and that the citizens found that the images added some value to their interpretation of the incident. Scores were also computed for the usability and efficiency of the interfaces, using post-condition questionnaires. Due to space limitations the results are not presented here. The interested reader can find all ER evaluation results in [79]. Overall both groups of participants were positive about the ER demonstrator application. They were able to complete the tasks required of them both in terms of uploading information, accessing information and also interpreting and interacting with the information in order to build an understanding of what incidents were happening in the city.

8.1.2 Evaluation of Consumer Social Group demonstrator

Regarding to the travel desktop prototype, the evaluation process was carried out with two different groups. One group, consisted by some of the WeKnowIt consortium members, had prior knowledge of the Collective Intelligence approach taken in the demonstrator and served as a provider of information that can be used in future enhancements of the prototype. The second group, on the other hand, were completely external users who showed their opinions for a product and not for a research prototype. Their opinion is especially important to observe how a tool like the travel desktop demonstrator would impact as a real product for planning a trip, aiming to help users to get the most complete view of places they want to visit.

For evaluation purposes, a generic Survey platform was implemented with the aim to provide an easy way to perform the evaluation from remote locations. The demonstrator was evaluated in terms of usability and satisfaction, while implicit user click feedback was collected to observe the usage of the tool and thus improve and refine workflows. Moreover, the effectiveness of its Collective Intelligence services was evaluated through qualitative questionnaires.

The SUS Usability and Satisfaction questionnaire was used, which allows for the usage of a standard methodology, which outputs a score on a scale from 0 to 100, the greater the score the more usability and satisfaction. The average SUS score obtained was 68.92 over 100 for the experienced group and 60.66 over 100 for external users, which is interpreted as a decent overall usability score with room for future improvement. The difference in the scores for the two user groups is not statistically significant ($P > 0.85$ using two-tailed unpaired t-test).




Fig. 8.6: Desktop travel application- SUS Average for both groups. Normalized values per question.

The interested reader can find more detailed results of the evaluation in [79]. In terms of using the evaluation to assess the quality of the Collective Intelligence services, it can be concluded that the search functionalities of the services (exploiting content from different sources) and the POIs clustering features were of highest importance to the end users.

Concerning the mobile part of the travel scenario demonstrator, two field test evaluations were performed, each implemented by different groups of evaluators. The first evaluation was performed in Madrid by a group of four people who perform a two-hour trip in Madrid. The second evaluation has been performed in Barcelona by a group of 21 persons, who evaluated the demonstrator during a 5-hours field test. The Mobile Guidance demonstrator suggested to the users different routes, according to their preferences. The user feedback was collected by means of a questionnaire. The evaluation addressed the following user needs: touristic information and recommendations, field navigation guidance and group communication, which were supported by Collective Intelligence automatically extracted results. The items of functionality considered for evaluation were: get recommendations of Points of interest; search for places; search for Points of Interest; search for friends; get detailed information about places; get detailed information about Points of interests; get image gallery about Points of interest.

The following dimensions were explored in the evaluation: availability, usefulness, usability, simplicity, appearance, responsiveness, suitability, completion and contribution of Collective Intelligence to the user experience.

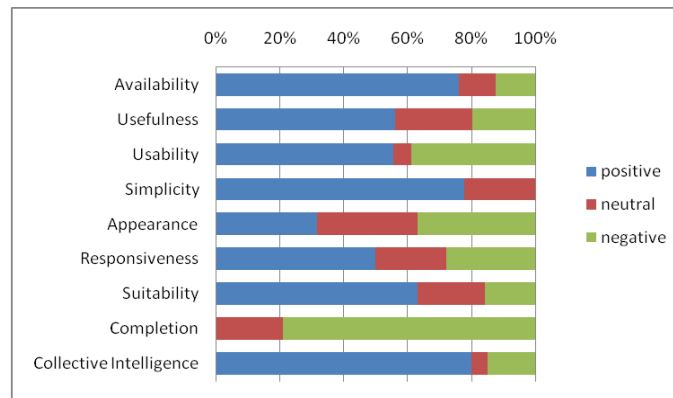


Fig. 8.7: Mobile guidance evaluation outcome.

The evaluation outcome is displayed in Fig. 8.7. Most dimensions received positive feedback, which means complete or partial agreement on that dimension. Most evaluators were positive on the ability to exploit Collective Intelligence in the travel scenario using the Mobile Guidance application.

9. Conclusions

We have presented state-of-the-art technologies from the media, mass, social and organisational intelligence layers which exploit massive Web 2.0 and user generated content. The methods can be used either solely or in combination, so as to provide an enriched level of intelligence, the so called Collective Intelligence.

The integration of the different technology layers which analyse large amount of user generated data in complementary approaches has led to the leveraging of an integrated Collective Intelligence framework, a software platform incorporating analysis services and functionalities from these diverse modalities.

Regarding the exploitation of single layer techniques, in media intelligence the developed VIRaL tool presents an integrated approach on visual image retrieval and localization. We show how the bag-of-words model can be extended by adding geometrical consistency into it and how geo-tags may be exploited in order to allow localization and POIs identification. The nature of this research and the hot topic it comprises within the Web 2.0 framework allows us to identify numerous extensions among the future goals of this work. The most important ones are the expansion of the utilized datasets to larger sets of publicly available, user-generated images, improvements on the algorithmic geometrical consistency model, extension of the presented content matching algorithm in order for it to be used for fast and robust processing of mass amount of non-geotagged images and additional integration of external social web sources as ad-hoc services.

In the mass intelligence layer, methods developed for ontology-based classification of documents create the basis for ontological navigation within thematically related documents. Analysis of semantic graphs created from the input documents can reveal overlapping information that constitutes the core of a topic, and complementary information covered only in a document subset. Such distinction between major and minor topics covered by document set will facilitate more fine-grained topical navigation within the given set of related documents. This approach can be used to create a comprehensive description of a topic and facets to efficiently navigate within it.

The developed community detection methods are a valuable tool for studying the mesoscopic structure of graph-based data, e.g. folksonomies, which represent associations among users, content items and tags. The study of the derived communities can be valuable for a series of tasks such as content indexing, tag clustering and tag recommendation. The efficiency of the developed methods make them suitable for tackling large-scale analysis problems.

In the context of the combination of the different intelligence layers, new Collective Intelligence techniques can be created that leverage the information of several sources as well as services that combine technologies from different research areas. For example, image similarities can be used to weight the networks of their photographers. New opportunities for social network analysis and recommendation systems arise.

On the organisational intelligence side, the ontology support for collaboratively creating and sharing semantic POIs can be transferred and used as modelling basis to support collaborative user activities in other applications and domains. Due to its domain independent design, the Event-Model-F can be applied in various other domains. First examples of using the Event-Model-F for an application in the domain of history of art and research of scientific art pieces have shown that it can be easily

specialized towards domain-specific requirements. The events represented using the Event-Model-F can also be associated with documentary support from the Web 2.0 such as images and Flickr and tweets on Twitter. Finally, one can conduct reasoning on the Event-Model-F by leveraging domain-specific knowledge such as in emergency response.

From the integration point of view, the produced Data Storage component for the ICIF is a generic solution suitable for storing data of different applications. It can be tailored to store business objects of any application. The WKI DS can be used to enrich typical business applications with SPARQL querying capabilities.

Apart from the possibilities on building on top of the results of each intelligence layer, the most important outlook of the presented Collective Intelligence approach comes from the fact that the aggregated benefits of Collective Intelligence acquired through the various Intelligence Layers are, in particular, realised by both the end users and the organisations when uploading, searching for, browsing and consuming the content. Therefore, each combination of technologies does not only contribute to the generation of Collective Intelligence resulting from each technique separately, but the integration of these varied techniques results in an overall added value which exceeds the aggregate of the individual techniques benefits.

Acknowledgments

The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement n°215453 - WeKnowIt.

References

- [1] Work, D, Blandin, S, Piccoli, B, Bayen, A (2010). A traffic model for velocity data assimilation. *Applied Mathematics Research eXpress*.
- [2] Torres, LH (2007). Citizen sourcing in the public interest. *Knowledge Management for Development Journal*, 3(1): 134-145.
- [3] iSpot, your place to share nature. URL: <http://ispot.org.uk/>.
- [4] Mobnotes. URL: <http://www.mobnotes.com/>.
- [5] Dopplr. URL: <http://www.dopplr.com/>.
- [6] FixMyStreet. URL: <http://www.fixmystreet.com/>.
- [7] Malone, TW, Klein, M (2007). Harnessing Collective Intelligence to Address Global Climate Change. *Innovations*, 2(3):15-26.
- [8] Kemp, C, Shafto, P, Berke, A, Tenenbaum, JB. Combining causal and similarity-based reasoning. *Advances in Neural Information Processing Systems* 19 , in press.
- [9] Prelec, D (2004). A Bayesian Truth Serum for Subjective Data. *Science*, October 2004, 462-466.
- [10] MIT Center for Collective Intelligence, Distributed collaboration project, URL: <http://cci.mit.edu/research/collaboration.html>.
- [11] Wang, C, Jing, F, Zhang, L, Zhang, H (2007). Content-based image annotation refinement. In *CVPR*.
- [12] Liu, D, Hua, X, Wang, M, Zhang, H (2010). Retagging social images based on visual and semantic consistency. In *WWW*.
- [13] Li, J, Wang, JZ (2003). Automatic linguistic indexing of pictures by a statistical modeling approach. *TPAMI*.
- [14] Chang, EY, Goh, K, Sychay, G, Wu, G (2003). Cbsa: content-based soft annotation for multimodal image retrieval using bayes point machines. *TCSVT*.
- [15] Jeon, J, Lavrenko, V, Manmatha, R (2003). Automatic image annotation and retrieval using cross-media relevance models. In *SIGIR*.
- [16] Feng, S, Manmatha, R, Lavrenko, V (2004). Multiple Bernoulli relevance models for image and video annotation. In *CVPR*.
- [17] Zhou, D, Bousquet, O, Lal, T, Weston, J, Scholkopf, B (2003). Learning with local and global consistency. In *NIPS*.
- [18] Kennedy, LS, Naaman, M, Ahern, S, Nair, R, Rattenbury, T (2007). How Flickr helps us make sense of the world: context and content in community-contributed media collections. *ACM Multimedia*, 631-640.
- [19] Quack, T, Leibe, B, Van Gool, L (2008). World-scale mining of objects and events from community photo collections. In *Proceedings of the 2008 international conference on Content-based image and video retrieval*, 47-56.
- [20] Girardin, F, Calabrese, F, Dal Fiore, F, Ratti, C, and Blat, J (2008). Digital footprinting: Uncovering tourists with user-generated content. *IEEE Pervasive Computing*, 7(4):36-43
- [21] Hollenstein, L, Purves, RS (2009). Exploring place through user-generated content: using Flickr to describe city cores. *Journal of Spatial Information Science*.
- [22] Lowe, DG (2004). Distinctive image features from scale-invariant keypoints. *International Journal of Computer Vision*, 60(2):91-110.
- [23] Nister, D, Stewenius, H (2006). Scalable recognition with a vocabulary tree. In *ACM Computer Vision and Pattern Recognition*, 27: 2161-2168.
- [24] Sivic, J, Zisserman, A (2003). Video google: A text retrieval approach to object matching in videos. In *International Conference on Computer Vision*, 2:1470-1477.
- [25] Chum, O, Philbin, J, Zisserman, A (2008). Near duplicate image detection: Min-hash and tf-idf weighting. In *ACM British Machine Vision Conference*, 2:1.

- [26] Grauman, K, Darrell, K (2007). Pyramid match hashing: Sub-linear time indexing over partial correspondences. In *Computer Vision and Pattern Recognition*, 1–8.
- [27] Philbin, J, Chum, O, Isard, M, Sivic, J, Zisserman, A (2007). Object retrieval with large vocabularies and fast spatial matching. In *IEEE Computer Society Computer Vision and Pattern Recognition*, 1:47–56.
- [28] Quack, T, Leibe, B, Van Gool, L (2008). World-scale mining of objects and events from community photo collections. In *ACM Content-based Image and Video Retrieval*, 2: 47–56.
- [29] Fischler, MA, Bolles, RC (1981). Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications*, 24(6):381–395.
- [30] Wolfson, HJ, Rigoutsos, I (1997). I. Geometric hashing: An overview. *IEEE Computational Science and Engineering*, 82(2):10–21.
- [31] Giannakidou, E, Koutsonikola, VA, Vakali, A, Kompatsiaris, Y (2008). Co-Clustering Tags and Social Data Sources. *Proceedings of WAIM 2008: 9th International Conference on Web-Age Information Management*. IEEE, 317–324
- [32] Brooks, CH, Montanez, N (2006). Improved annotation of the blogosphere via autotagging and hierarchical clustering. *Proceedings of WWW '06: 15th international Conference on World Wide Web*. ACM, New York, NY, 625–632.
- [33] Begelman, G, Keller, P, Smadja, F (2006). Automated Tag Clustering: Improving search and exploration in the tag space. Online article: http://www.pui.ch/phred/automated_tag_clustering.
- [34] Vapnik, V (1995). *The nature of statistical learning theory*. Springer Verlag.
- [35] Lewis, DD (1998). Naive (Bayes) at forty: The independence assumption in information retrieval. *ECML-98, 10th European Conference on Machine Learning*, Chemnitz, DE.
- [36] Deerwester, S, Dumais, S, Furnas, GW, Landauer, TK, Harshman, R (1990). Indexing by Latent Semantic Analysis. *Journal of the Society for Information Science*, 41:391–407.
- [37] Bloehdorn, S, Hotho, A (2004). Text Classification by Boosting Weak Learners based on Terms and Concepts. *4th IEEE International Conference on Data Mining (ICDM'04)*.
- [38] Gabrilovich, E, Markovitch, S (2006). Overcoming the Brittleness Bottleneck using Wikipedia: Enhancing Text Categorization with Encyclopedic Knowledge. *21th National Conference on Artificial Intelligence*, Boston, MA, USA.
- [39] Winerman, L (2009). Social networking: Crisis communication. *Nature*. 457(7228):376.
- [40] Rissanen, E (2010). eXtensible Access Control Markup Language (XACML) Version 3.0 Committee Draft 03. In *Parducci, B, Lockhart, H (eds.), OASIS eXtensible Access Control Markup Language (XACML) TC*, URL: <http://docs.oasis-open.org/xacml/3.0/xacml-3.0-core-spec-cd-03-en.pdf>.
- [41] Anderson, AH (2006). A comparison of two privacy policy languages: EPAL and XACML. In *Proceedings of the 3rd ACM workshop on Secure web services*, Alexandria, Virginia, USA, 53–60.
- [42] Raghavan, UN, Albert, R, Kumara, S (2007). Near linear time algorithm to detect community structures in large-scale networks. *Phys. Rev. E* 76, 036106.
- [43] Wang, X, Mamadgi, S, Thekdi, A, Kelliher, A, Sundaram, H (2007). Eventory -- An Event Based Media Repository. *Semantic Computing*, IEEE.
- [44] Raimond, Y, Abdallah, S (2004). The Event Ontology, October, 2007. URL: <http://motools.sf.net/event>
- [45] Doerr, M, Ore, CE, Stead, S (2004). The CIDOC conceptual reference model: a new standard for knowledge sharing, Conceptual modelling. *Australian Computer Society, Inc.*
- [46] IPTC, EventML, 2008. URL: <http://iptc.org/>
- [47] Mueller, ET (2008). Event Calculus. *Handbook of Knowledge Representation*, Elsevier.
- [48] Ekin, A, Tekalp, AM, Mehrotra, R (2004). Integrated semantic-syntactic video modeling for search and browsing. *IEEE Transactions on Multimedia*, 6(6).
- [49] Francois, ARJ, Nevatia, R, Hobbs, J, Bolles, RC (2005). VERL: An Ontology Framework for Representing and Annotating Video Events. *IEEE MultiMedia*, 12(4).
- [50] Westermann, U, Jain, R (2007). Toward a Common Event Model for Multimedia Applications. *IEEE MultiMedia*, 14(1).
- [51] Scherp, A, Franz, T, Saathoff, C, Staab, S (2009). F-A Model of Events based on the Foundational Ontology DOLCE+ Ultra Light. *K-CAP*.
- [52] Aims: Atlas incident management system. URL: <http://www.atlasops.com/products/aims.php>.
- [53] Emergency command system. URL: <http://www.emergencycommandsystem.com>.
- [54] Using geography can help you to meet your flood management responsibilities. URL: <http://www.ordnancesurvey.co.uk/oswebsite/business/sectors/government/local/news/flood-management.html?cmp=floodmgt>
- [55] MacQueen, J (1967). Some methods for classification and analysis of multivariate observations. In *Proceedings of 5th Berkeley Symposium on Mathematical Statistics and Probability*, 1–297.
- [56] Jegou, H, Douze, M, Schmid, C (2008). Hamming embedding and weak geometric consistency for large scale image search. In *Andrew Zisserman David Forsyth, Philip Torr, editor, European Conference on Computer Vision, LNCS*. Springer.
- [57] Kalantidis, Y, Tolias, G, Spyrou, E, Mylonas, P, Avrithis, Y (2009). Visual Image Retrieval and Localization. ViRaL demo presentation at the 7th International Workshop on Content-Based Multimedia Indexing (CBMI 2009), Chania, Greece.
- [58] Mika, P (2005). Ontologies are us: A unified model of social networks and semantics. *ISWC 2005, LNCS 3729*, 522–536, Springer-Verlag.
- [59] Papadopoulos, S, Kompatsiaris, Y, Vakali, A (2010). A Graph-based Clustering Scheme for Identifying Related Tags in Folksonomies. In *DaWaK 2010: the 12th International Conference on Data Warehousing and Knowledge Discovery*.
- [60] Xu, X, Yuruk, N, Feng, Z, Schweiger, TA (2007). SCAN: A Structural Clustering Algorithm for Networks. *Proceedings of KDD '07: 13th international Conference on Knowledge Discovery and Data Mining*, ACM, New York, NY, 824–833.
- [61] Luo, F, Wang, JZ, Promislow, E (2006). Exploring Local Community Structures in Large Networks. *Proceedings of the 2006 IEEE/WIC/ACM international Conference on Web Intelligence*. IEEE Computer Society, 233–239.

- [62] Papadopoulos, S, Vakali, A, Kompatsiaris, Y (2010). Community Detection in Collaborative Tagging Systems. In Book Community-built Database: Research and Development, Springer.
- [63] Clauset, A, Newman, MEJ, Moore, C (2004). Finding community structure in very large networks. *Physical Review E* 70, 066111.
- [64] Janik, M, Kochut KJ (2008). Training-less ontology-based text categorization. In Workshop on Exploiting Semantic Annotations in Information Retrieval (ESAIR 2008) at the 30th European Conference on Information Retrieval (ECIR'08), Glasgow, Scotland.
- [65] Janik, M, Kochut KJ (2008). Wikipedia in action: Ontological knowledge in text categorization. In 2nd International Conference on Semantic Computing (ICSC), Santa Clara, CA, USA.
- [66] Kleinberg, JM (1998). Authoritative sources in a hyperlinked environment. In ACM-SIAM Symposium on Discrete Algorithms.
- [67] Lewis, DD, Yang, Y, Rose, T, Li, F (2004). Rcv1: A new benchmark collection for text categorization research. *Journal of Machine Learning Research*, 5:361–369.
- [68] McCallum, AK (1996). Bow: A toolkit for statistical language modeling, text retrieval, classification and clustering.
- [69] Witten, IH, Frank, E (2005). *Data Mining: Practical machine learning tools and techniques* (2nd ed.). Morgan Kaufmann, San Francisco.
- [70] Ovelgönne, M, Geyer-Schulz, A, Stein, M (2010). A randomized greedy modularity clustering algorithm for community detection in huge social networks. 4th international Workshop on Social network Analysis and Mining (SNA-KDD 2010), Washington DC, USA.
- [71] Newman, MEJ, Girvan, M (2004). Finding and evaluating community structure in networks. *Physical Review E*, 69(2).
- [72] Franke, M, Geyer-Schulz, A (2009). An update algorithm for restricted random walk clustering for dynamic data sets. *Advances in Data Analysis and Classification*, 3(1):63-92.
- [73] Sandhu, RS, Samarati, P (1994). Access control: principle and practice. *Communications Magazine, IEEE*, 32(9): 40-48.
- [74] Sikkil, K (1997). A Group-based Authorization Model for Cooperative Systems. In *Proceedings of the Fifth European Conference on Computer-Supported Cooperative Work*, Springer.
- [75] Geyer-Schulz, A, Thede, A (2006). Implementation of hierarchical authorization for a web based digital library. In 3rd International Conference on cybernetics and information technologies, systems, and applications, 139-144.
- [76] Ferraiolo, DF, Kuhn, DR (1992). Role-Based Access Controls. In *Proceedings of the 15th National Computer Security Conference*, Baltimore MD, 554 – 563.
- [77] Saathoff, C, Scherp, A (2010). Unlocking the semantics of multimedia presentations in the web with the multimedia metadata ontology. *WWW conference*.
- [78] 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).
- [79] WeKnowIt project deliverable D7.5.1: Consumer and emergency response use case first evaluation report. URL: <http://www.weknowit.eu/deliverables>
- [80] Kalantidis, Y, Tolias, G, Avrithis, Y, Phiniketos, M, Spyrou, E, Mylonas, Ph, Kollias, S (2011). VIRaL: Visual Image Retrieval and Localization. *Multimedia Tools and Applications*, (to appear).
- [81] Avrithis, Y, Kalantidis, Y, Tolias, G, Spyrou, E (2010). Retrieving Landmark and Non-Landmark Images from Community Photo Collections. In *Proceedings of ACM Multimedia (MM 2010)*, Firenze, Italy.
- [82] Papadopoulos S, Zigkolis C, Tolias G, Kalantidis T, Mylonas P, Kompatsiaris Y, Vakali A (2010). Image clustering through community detection on hybrid image similarity graphs. In *Proceedings of ICIP 2010: International Conference on Image Processing*. IEEE.
- [83] Geyer-Schulz, A, Ovelgönne, M, Sonnenbichler, AC (2010). Getting Help In A Crowd - A Social Emergency Alert Service. In *Proceedings of the International Conference on e-Business 2010 (ICETE ICE-B)*, Athens, Greece, 207-218.
- [84] Ovelgönne, M, Sonnenbichler, AC, Geyer-Schulz, A (2010). Social Emergency Alert Service - A Location-Based Privacy-Aware Personal Safety Service. In *Next Generation Mobile Applications, Services and Technologies (NGMAST)*, 84-89.
- [85] Baader, F, Calvanese, D, McGuinness, DL, Nardi, D, Patel-Schneider, PF, editors (2003). *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press. ISBN 0-521-78176-0.
- [86] OWL 2 Web Ontology Language. URL: <http://www.w3.org/TR/owl2-overview/>
- [87] The Protégé Ontology Editor and Knowledge Acquisition System. <http://protege.stanford.edu/>