WeKnowIt

Emerging, Collective Intelligence for Personal, Organisational and Social Use

FP7-215453

D4.4

Implementation of Service Modules for Cross-Usage

Dissemination level	Confidential
Contractual date of delivery	31.12.10
Actual date of delivery	30.12.10
Work package	WP4 Social Intelligence
Task	T4.1 Community Service exploiting the Cross-Usage of Intelligence
Type	Prototype
Approval Status	Approved
Version	13
Number of pages	59
Filename	D4.4-man_2011-04-02_v14_CERTH_deliverable.odt

Abstract

The Social Layer of the WeKnowIt System integrated three new and explicit cross-usage services. The Emergency Alert Service is a location-based service that activates nearby members of the social group of a victim in the case of an emergency. The reputation service formalizes the assessment of user generated content in the WeKnowIt system. The recommendation service aims to exploit the natural multi-modality of social media resources uploaded on the WKI platform by users, analyze and combine intelligence from different modalities and offer recommendations that are based on a joint weighted use of *Media*, *Mass* and *Social Intelligence.*

SEVENTH FRAMEWORK PROGRAMME

co-funded by the European Union

History

Version	Date	Reason	Revised by
1	22.10.10	Creation / Structure	Andreas Sonnenbichler
2	11/10/10	Added Section	Eirini Giannakidou
3	29/11/10	Added Section	Andreas Sonnenbichler
4	13/12/10	Added Section	Felix Schwagereit
5	13/12/10	Fused sections	Andreas Sonnenbichler
6	13/12/10	Added header sections	Andreas Sonnenbichler
7	13/12/10	Further section processing	Eirini Giannakidou
11	30/12/10	Finalizing	Andreas Sonnenbichler

Author list

Organization	Name	Contact Information
EMKA	Andreas Sonnenbichler	Karlsruhe Institute of Technology
EMKA	Michael Ovelgönne	Karlsruhe Institute of Technology
UoKob	Felix Schwagereit	University of Koblenz
CERTH	Eirini Giannakidou	Informatics and Telematics Institute

Executive Summary

This deliverable D4.4 "Implementation of service modules for cross-usage belongs" to task T4.1 "Community services exploiting the cross-usage of intelligence ". It aims to describe services developed in this task making use of the layers of collective intelligence which are hosted and linked to the WP4 Social Intelligence layer.

In this document we will provide three services:

The analysis of social networks focuses on scalable algorithms for community detection in networks. This capability of inferring social groups from huge communication networks is important for the **Emergency Alert Service** (EAS) – a prototype application for WeKnowIt's scenarios. The EAS is a location-based service that activates nearby members of the social group of a victim in the case of an emergency. It is a mobile service which can be installed as an applet on smart phones (e.g. Android based phones). When activated it works as an emergency call agent and informs social contacts (friends, family, and colleagues) and public authorities about the emergency situation. Current geo-position and routing information is provided. The service is designed privacy-aware and dynamic. It is useful for the Consumer Group and the Emergency Response Scenario.

A key aspect in the **Recommendation Service** (RecS) is to enable the user leverage knowledge from the different kinds of intelligence generated in the WKI context so far. To this end, the service aims to exploit the natural multi-modality of social media resources uploaded on the WKI platform by users, analyze and combine intelligence from different modalities and offer recommendations that are based on a joint weighted use of *Media*, *Mass* and *Social Intelligence.* The outcome of the service is expected to benefit the Consumer Group Scenario.

The **Reputation Service** (RepS) describes a way to formalize the assessment of user generated content through explicit reputation. We consider reputation as a quantitative value that explicitly represents a property of an entity of interest, like users and content. The purpose of this reputation system is to assign a reputation to these entities that has a correlation with a non-visible property. It will be used especially in the ER use case scenario for collecting and analysing the reliability of user-generated content. Since this content can be the basis for important decisions made by ER personnel quality considerations apply especially in this case. The reputation system therefore provides its reliability scores for users and content items based on intelligence provided by work packages WP1, WP3, WP4, and WP5.

Abbreviations and Acronyms

AC	Access Condition
ACL	Access Control List
AJAX	Asynchronous JavaScript and XML
BNF	Backus-Naur-Form
CAP	Community Administration Platform
CAT	Community Analysis Tool
CDL	Community Design Language
CMLM	Community Membership Life Cycle Model
CMS	Content Management System
EC	Emergency Case
ECL	Emergency Case Leader
ECM	Emergency Case Member
RBAC	Role Based Access Control
RDF	Resource Description Framework
SNA	Social Network Analysis
SPARQL	SPARQL Protocol and RDF Query Language
URI	Uniform Resource Identifier
WKI	WeKnowIt
WP	Work Package
XML	eXtensible Markup Language
EAS	Emergency Alert Service
RepS	Reputation Service
RecS	Recommendation Service

Table of Contents

Illustration Index

Index of Tables

1. The Emergency Alert Service

This chapter describes the Emergency Alert Service developed with work package 4 Social Intelligence Layer.

1.1. Introduction

On Saturday, the 12th of September 2009, 50-year old Dominik Brunner was brutally murdered in a Munich S-Bahn station [1]. The attack on Dominik Brunner was observed by 15 passengers [2] and transmitted and recorded by his mobile-phone on the open police emergency channel [3].

This tragic incident reminds social psychologists of the murder of Kitty Genovese on March 13th, 3:20, 1964 in Queens, New York. She was stalked, stabbed, and sexually assaulted near her own apartment building. During the attack on her she screamed and broke free twice. 38 of her neighbours witnessed the attack, but no one intervened. After 45 minutes one man called the police, but at this point in time Kitty Genovese was dead.

This incident motivated the study of social processes in emergency situations by Darley and Latane [4] and it points to the short-comings and problems of real emergency response organizations and their management which only very recently have become the object of scientific research e.g. [5], [6], [7], and [8].

The mobile phone of the victim transmitting and recording to the end confirms Palen and Liu's thesis of the increasing availability of ICT and its use in an emergency by citizens and also their observation, that the traditional hierarchical command-and-control reporting system of emergency response organizations may not be adequate and "does not include built-in considerations for the important roles that members of the public play as participants" [9, p. 729]. Public participation in emergencies and disasters is active and altruistic. First responders are often not the trained professionals of an emergency response organization who are sent to the incident, but the people from the local and surrounding communities. They provide first-aid, transport victims to the hospital, and begin search and rescue [9, pp. 728-729].

This chapter starts with a short review of the social processes which Darley and Latene have identified as obstacles for helping in emergencies in chapter 1.2.

These obstacles have become known as the bystander effect: The more bystanders, the less likely the victim will receive help. We address the bystander effect by a social emergency alert service and discuss how social emergency alert services may help in improving these processes in emergency settings by activating the nearby members of the victim's social network.

For the process of giving and receiving help, we propose to monitor social interactions and to identify the social groups of the victims and to locate the nearest members of the social group of the victim in an emergency for notification purposes.

In chapter 1.3. we present details on the social emergency alert service for getting help in a crowd and in chapter 1.4. we discuss implementation variants of realizations of such services based on readily available technology by the telecommunication and Internet industry.

In chapter 1.5. we will discuss two system designs that emerged from our work on a prototype of the emergency alert service.

The identification of social groups from social networks will be addressed in chapter 1.6..

Chapter 1.7. aims at assessing the chances that the social emergency alert service presented has in reality. For this purpose, a first attempt is made to answer three questions which play a crucial role for the success of a social emergency alert service:

- Has the victim in an emergency a chance to transmit a request for help?
- Is someone of his social network nearby?
- Will this person really help?

1.2. The Unhelpful Crowd: Five steps to helping in an Emergency

The murder of Kitty Genovese in 1964 in Queens, New York, in front of 38 witnesses who did not interfere led Darley and Latane to start investigating the social psychological processes at work in this incident. Their research revealed that the more bystanders, the less likely the victim will be helped. This is the **bystander effect**: The presence of others inhibits helping.

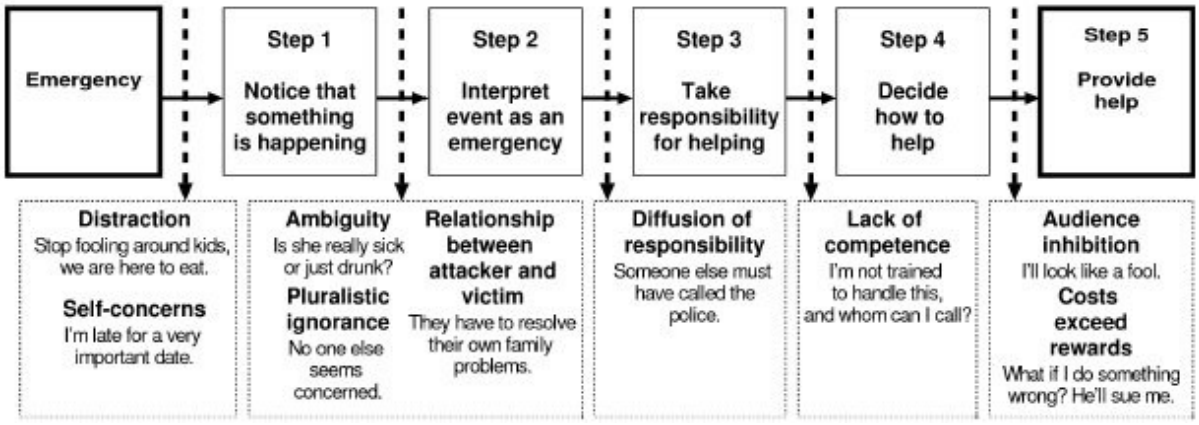

Figure 1: The social help process: 5 steps leading to assistance.

One of the reasons for this is the diffusion of responsibility [10]. Darley and Latane [4] provided a careful analysis of the process of emergency intervention shown in figure 1. Unfortunately, many obstacles to helping must be overcome. At each step, however, psychological factors are at work which explain why people fail to help:

1. Noticing
   The presence of others distracts attention from the victim. People who live in big cities may filter out people lying on sidewalks or screams (stimulus overload [11]).
2. Interpreting
   People must interpret the meaning of what they perceive. Their perception may be ambiguous: Cries of pain may be taken for laughter, hypoglycemia may be taken for drunkenness (e.g. [12] [13]).
   A perceived relationship between attacker and victim may lead the observers to think that everything is OK [14]. If an emergency happens, the most powerful information available is often what other people do. However, if everybody looks on everybody else to get clues on what to do, the whole group is suffering from pluralistic ignorance and everybody concludes that help is not needed (e.g. [15] and [16].
3. Take responsibility
   When help is needed, who is responsible for providing it? The diffusion of responsibility means that people belief that others will or should help. The effect usually is strengthened by anonymity and considerably reduced by a reduction in psychological distance. Groups with members who know each other are more helpful than strangers. See e.g. [17], [18], and [19].
4. Decide how to help
   Bystanders are more likely to offer direct help if they feel competent to perform the actions required (e.g. [20], [21]).
5. Provide help
   Some people may feel too embarrassed to provide help in a public setting (audience inhibition). However, when people think they will be scorned by others for failing to provide help, the presence of an audience will increase their helpful actions. See [22].
   In addition, potential helpers - especially when uncertain about their capability to help - tend to weight the risks of helping against the reward of helping others - because of costs exceeds rewards - remain passive.

In addition, a series of other variables have a high influence on helping behaviour as experiments in social psychology have shown: Time pressure reduces the tendency to help (e.g. [23], [24]). Group membership and empathy and attractiveness interact: empathy is a positive predictor for help for in-group members, whereas attractiveness works for out-group

members [25]. Group membership positively influences help for in-group members, and the group boundaries can be shifted by proper priming [26]. Group status and group identification influence the willingness of receiving help [27].

But what can you do to receive help in a crowd? Try to counteract the ambiguity of the situation by making it clear that you need help, and try to reduce the diffusion of responsibility by addressing a specific individual for help, keep eye contact, point or direct requests (e.g. [28], [29]).

Consistent with this is a recent study of P. Markey [30] of people in Internet chat rooms: If the number of individuals is large in a chat room, individuals react slower to a plea for help. However, addressing a specific individual by his name leads to considerably faster help and eliminates this effect.

Research on the bystander effect in social psychology showed that even weak social links matter and increase the chance of a victim to receive help considerably. This fact is the main motivation to send alerts to the geographically close members of the victim's social group.

The asymmetric perception of social links (e.g. [31]), the role of weak ties, and the cultural norms of the community play a major role in the formation of the social group of the victim. The asymmetric perception of social links implies that a person may not be really aware of possible helpers in his loose social contacts. Taken together with cultural norms, even professional acquaintances are potential helpers.

The role of weak ties for networks has been studied by M. Granovetter ([32], [33]). In the context of information diffusion on open jobs Granovetter observed that "it is remarkable that people receive crucial information from individuals whose very existence they have forgotten" [33, p. 14]. This is an indication that an explicit list of emergency contacts provided by the subscriber of such an emergency service will considerably limit the effectiveness of such a service, because of these social phenomena.

1.3. A Social Emergency Alert Service

As a consequence of the problem of getting help in a crowd we propose a social emergency alert service (EAS) that identifies nearby members of the social group of the victim and notifies them about the victim's need for help and the victim's location. With this service we aim to activate the locally available social network of the victim and to eliminate the bystander effect.

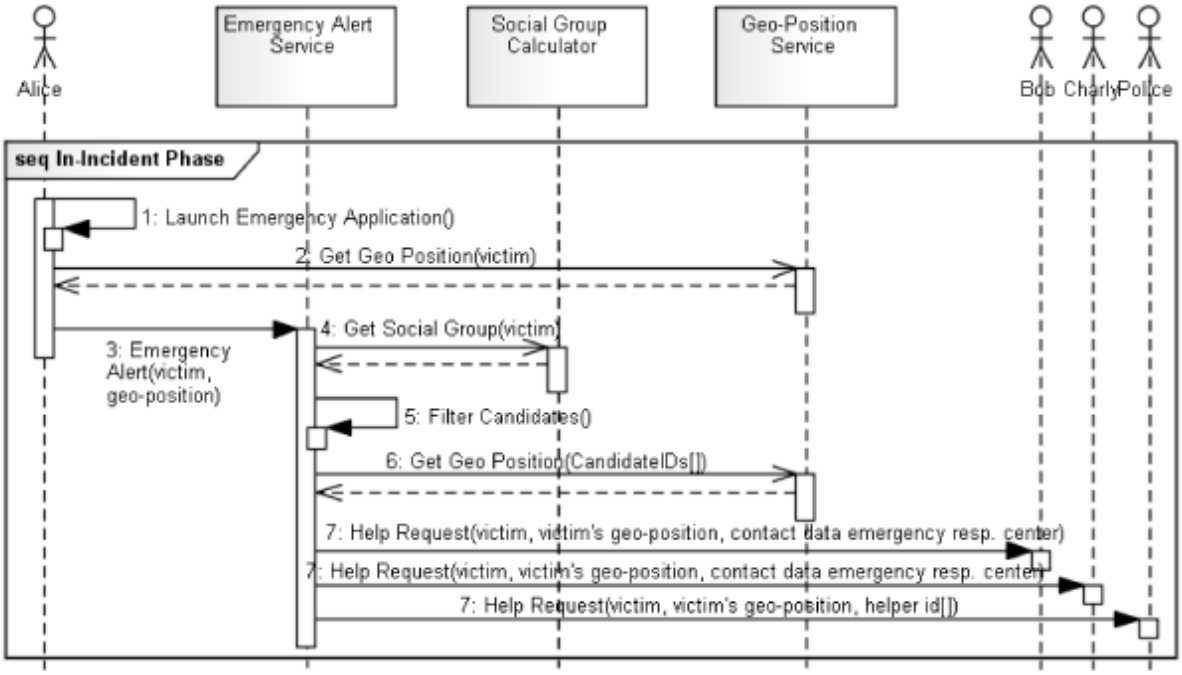

Figure 2: Sequence diagram of a generic social emergency alert service (in-incident phase)

The UML sequence diagram depicted in figure 2 shows the generic process in an emergency incident. It is designed on a high-level, abstract way allowing a variety of industrial implementations. We will address this issue in chapter 1.4.

An emergency notification is submitted by the victim by starting an application on his mobile device (`LaunchEmergenyApplication` in figure 2) e.g. by pressing the help-button shown in figure 3. The application retrieves the current geo-position. Both, the ID of the emergency caller and his geo-location are then transmitted to the emergency alert service (`Emergency Alert Service` in figure 2).

Figure 3: Alert widget on an Android homescreen

Figure 4: Notification shown to helpers with own position, position of the victim and the shortest route

For discovering the most likely person to help in the victim's social network, his social network has to be known and possible helpers identified (`Social Group Calculator` in figure 2). The network is either built from social interaction data from e-mail, sms, phone, and mobiles where the number of interactions is taken as an indication of social nearness or from social web sites as for example Facebook or Xing.

The user has not to build the network manually. Instead it is enough if he allows the service to track/fetch the data and calculate the network. The motivation is straight forward: If a user is willing to use the service we expect him to be willing to share data necessary to provide the necessary social network data to find out about his contacts expected to help in his own emergency situation.

However, the number of social interactions may be ambiguous as a recent incident 34 of a woman threatened by her former husband with Googles Buzz has shown: So the possibility to check for such unwanted relations must be provided for the participants of the emergency alert service (`Filter Candidates` in figure 2).

Usually social networks tend to be very large. As emergencies are often time-critical, it might take too long to process such a network on-the-fly. Therefore, the network is generated and updated regularly for all service subscribers. To find out which persons in the social network are likely to help, a clustering of the network is performed. Details for this clustering are addressed in chapter 1.6.

Next, the current geo-position of the candidates is retrieved from a geo-position service (`GetGeoPositionofCandidates` in figure 2).

The alert service uses the geo-data as a filter on the victim's social group to find out, who of the possible helpers is locally close enough. Chapter 1.7. deals with details of having at least one member of the victim's social group in range.

The possible helpers in range and the emergency response center are informed about the emergency situation of the victim (several invocations of `HelpRequest` in figure 2). All possible helpers in range are informed simultaneously. The victim's name, his geo-location and the shortest route as well as the contact data for the emergency response center are provided.

Finally, possible helpers and the emergency response experts at the police's emergency response center may communicate, because of the information forwarded by the emergency alert service (not shown in figure 2).

This facility has the potential of providing expert guidance to the socially close first responders on the scene. However, it also reveals the identity of potential helpers to the emergency response center. The privacy impact of this must be addressed for such a service.

In chapter 1.2. obstacles to the five steps leading to assistance have been described. The emergency alert service presented addresses these obstacles directly:

1. The **distraction** obstacles can be avoided by noticing, that an emergency incident takes place: Clear signal words are part of the personal message to the helpers. This makes it obvious, an emergency case is happening and this is made clear to the helper.
2. **Self concerns** are also addressed by the service: Since the potential helper is directly addressed and others know this from the incident protocol, social and legal norms lead to pressure to help.
3. **Ambiguity** is by-passed by the clear and unmistakable help request sent to the helpers.
4. As this message is personal, **pluralistic ignorance** is eliminated. Experimental evidence for the elimination of pluralistic ignorance in internet chat rooms is provided by P. Markey [57].
5. The **relationship between attacker and victim** can not lead the helper to overlook the emergency event, because of the unambiguous emergency message.
6. **Diffusion of responsibility** is also reduced, since the emergency alert message is directly and personally addressed to the helper. Because of this, he is responsible and because of the incident protocol, others will know this and hold him responsible.
7. **Lack of competence** may be addressed by providing fast expert backup for helpers from the police emergency response center.
8. Expert guidance of how to help also addresses the problems of **audience inhibition** and **costs exceed rewards**.

1.4. Implementation Variants

In this chapter we present implementation examples how the **Emergency Alert Service** can be implemented in an industrial environment.

1.4.1. Emergency Alerting

To be able to use the service, the user has to possess a mobile device (e.g. a mobile phone). He can then subscribe to the service. In case of an emergency, he starts an application on his mobile device. Of course, the start of the application must be made simple and fast, as we do not expect it likely that the victim is able to deal with complex applications in an emergency situation. For the implementation third-party platforms like

Android can be used. Android[1] is a mostly free and open-source OS platform developed and driven by the Open Handset Alliance[2].

Further platforms like Apples iPhone may be supported as well. Special mobile devices combined with body-sensors, e.g. for elderly people, linking the start of the application to a hardwired button can be offered.

1.4.2. Geo-Position Service

Geo-positions of both the victim and all possible helpers of the victim's social network need to be calculated. Many of todays mobile phones are able to calculate their geo-position by GPS (Global Positioning System). The service Google Latitude is an example for a service that users can publish their current geo-position and share it with friends. If a mobile device does not include such a feature, several alternative techniques have been described and implemented. Even speed vectors can be calculated (for example [35] and [36]).

By this, the expected geo-position of somebody moving in a train can be found out.

1.4.3. Acquiring Social Network Data and Identifying Possible Helpers

For discovering the most likely person to help in an emergency case the social network has to be known. We present three possible realizations.

The social network can be built by monitoring outgoing and incoming calls on the mobile device of the subscriber. The emergency alert application running on the mobile device collects this call data, pools it and regularly (e.g. once a week) transmits it via HTTP to the social group service

(see figure 2). There the call logs of all service subscribers are combined and the communication network is generated: Telephone numbers are represented as nodes, the calls are weighted ties. Each call strengthens a tie.

The advantage of this solution is, that the network is independent from the telephone provider. It works depending just on the emergency alert application. The disadvantage is, that the calculated social network consists only of subscribers and their direct connections. Ties between non-subscribers can not be observed technically.

Alternatively not the mobile devices monitor the calls, but connection records from telephone providers are used. Connection records are stored for billing purposes. In the European Union an directive forces the provider to save call logs from six month up to two years [37].

These connection data can be used to calculate the social networks. Every connection is represented by a tie between the calling parties (more

1 Android: http://www.android.com/

2 Open Handset Alliance: http://www.openhandsetalliance.com/

concrete, their telephone numbers) as nodes. Of course, the resulting network will be huge. In chapter 1.6. we will discuss a clustering algorithm that is able to process such huge datasets. Other scalable algorithms exist (e.g. [38] and [39]).

The advantage of this solution is, that much more network data can be collected so that the problem of missing links is smaller. On the other hand this alternative can only be realized if the calling logs are available to the emergency alert provider. As we do not expect network providers to give such information away, the most likely approach for this alternative is, if the network provider is identical with the emergency alert service provider. The provider can then use the service as an additional opt-in feature. Another disadvantage is, that one network provider will probably not exchange network or call log data with other providers. By this the social network is limited to the customers of the provider plus their direct links.

As an additional feature of both alternatives address books in the customers' devices can be used to group telephone numbers. Different telephone numbers of one person can be combined and fused to one node in the social network.

A third approach to build the social network is to cooperate with existing online social network platforms. Data from Facebook, LinkedIn, Myspace can be used. The advantage of this solution is, that no subscriber or network provider boundaries exist. The disadvantage is, that people tend to accept more 'friends' in social platforms than they would accept offline. Additionally, most of these platforms do not weight their ties, so that the strength of a connection is unknown.

In practice all three alternatives used to build a social network as discussed above can be complemented with a list of emergency contacts provided by each subscriber and, if available, with a list of dedicated helpers for an event or for a community. In a German small rural community, the community's voluntary fire-fighters are an example of such a community. We expect, that people in the same social group are likely to help each other.

1.5. System Design

In chapter 1.4. critical parts of an emergency alert service have been discussed. Now, we turn to the system as a whole and discuss important issues of the system architecture and the communication protocols. We will present and discuss two approaches for the system architecture. The first approach is a centralized architecture where all system logic and data is provided by a server. The second approach is a less centralized and compared to the first approach some system logic and most user-location

data is shifted from the server to the mobile clients. A detailed presentation of this prototype can be found in [40].

To evaluate and improve the designs we implemented prototypes of them. Software components for the mobile clients have been developed for a

Motorola Milestone smartphone which is running the Android operating system. Server side components have been implemented as Servlets running within an Apache Tomcat environment.

1.5.1. Centralized Architecture

In figure 5 the system design for an alert service with a centralized architecture is shown (with the geo-position cache).

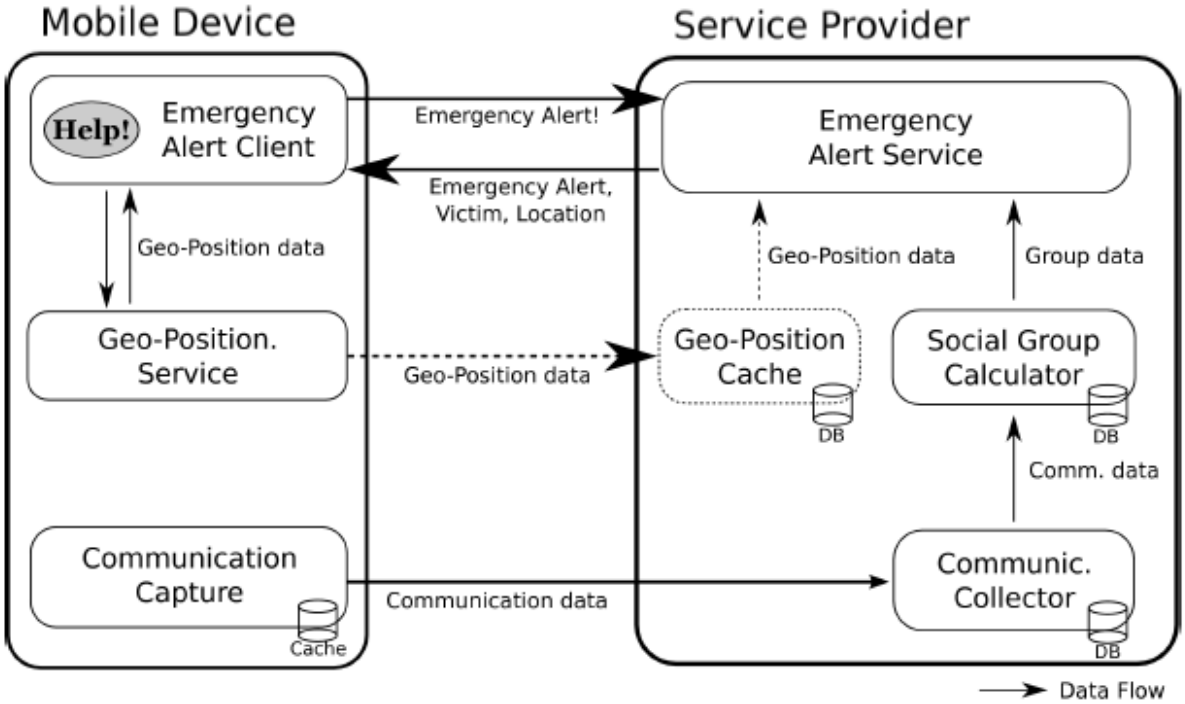

Figure 5: Architecture of the Emergency Alert Service. All boxes and arrows depicted in solid lines are components of the centralized and the privacy-aware system design. The Geo-Position Cache and its ingoing and outgoing data flow is depicted in dashed lines as this part becomes unnecessary in the privacy-aware system design.

Figure 6 depicts the corresponding communication protocol between the emergency alert service server components and the client software running on mobile devices.

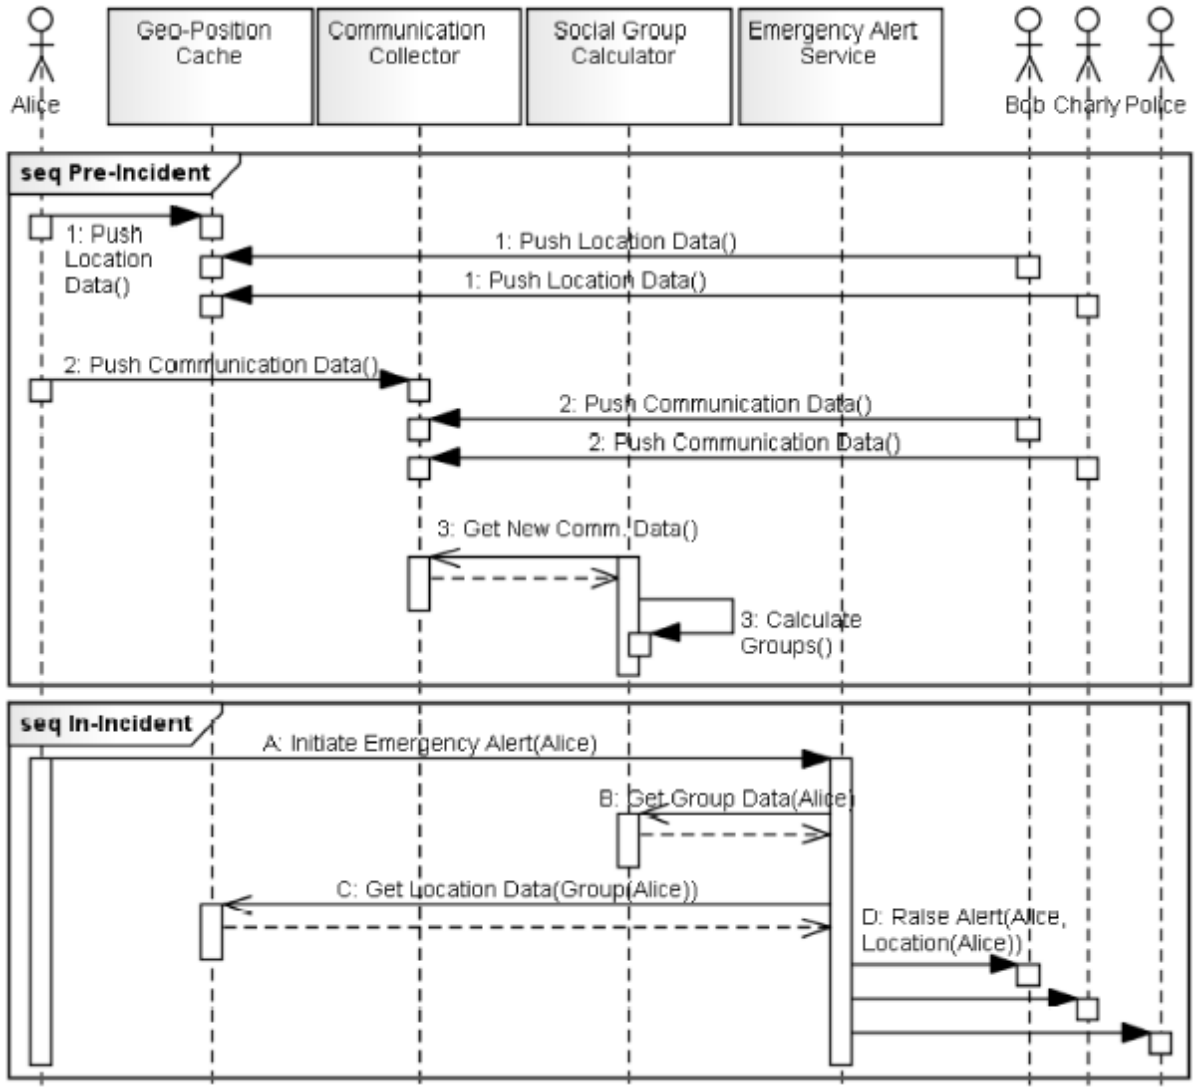

Figure 6: Communiction Protocol of a centralized emergency alert service.

Each subscriber is identified by his unique telephone number (ID). In the **Pre-Incidence Phase** each subscriber continuously submits location updates to the GPS Cache (step 1). We use a GPS Push Service running as background service on the Android smart phone to transmit the GPS coordinates via REST to GPS Cache. The transmission is done on a regular basis, e.g. every 10 minutes and if the position changed by more than 10 meters.

The GPS Cache stores the latest transmitted geo-position of each subscriber in a database. The second background service running on the mobile device is the Communication Capture Service. This service monitors all ingoing and outgoing calls and messages. The communication data is collected in a local cache. Once a day the smart phone submits the cached communication data via SOAP to the Communication Collector (step 2).

The Social Group Calculator pulls new data from the Communication Collector on a daily basis. It builds a communication network, identifies

social groups and stores the results in a database (step 3). OSGi is used as the internal Service Provider's protocol.

In the **In-Incident Phase** the user in need (in our example this is Alice) starts a home screen widget on her smart phone (see figure 3). A emergency alert is transmitted via REST to the Emergency Alert Service (step A). The service pulls the helpers candidate list from the Social GroupProvider for Alice via OSGi (step B). For each candidate the GPS Cache} is searched for the latest geo-position data of the candidate (step C). If close enough, a MMS is sent out to the helper with information about the victim and the victims geo-position (depicted on a map).

Note, that figure 6 is restricted to a proper emergency alert. False alerts can be revoked by a similar process (not shown in figure 6), which is password protected. However, pragmatically a set of passwords is provided which act as silent signals. A small solution consists of three passwords, the first signaling a false alarm, the second signaling that the victim is forced to revoke the alarm, and the third that there is danger for the helpers.

1.5.2. Privacy-Aware Architecture

The architecture presented in chapter 1.5.1. has two major drawbacks, both related to the continuous location transmission. The first is related to the energy consumption, the second to privacy concerns. See chapter 1.5.3.

To solve both issues we propose an alerting protocol that works without continuous location transmission to a server (see figure 7). This modification of the centralized approach requires to transfer some of the system logic to the mobile devices.

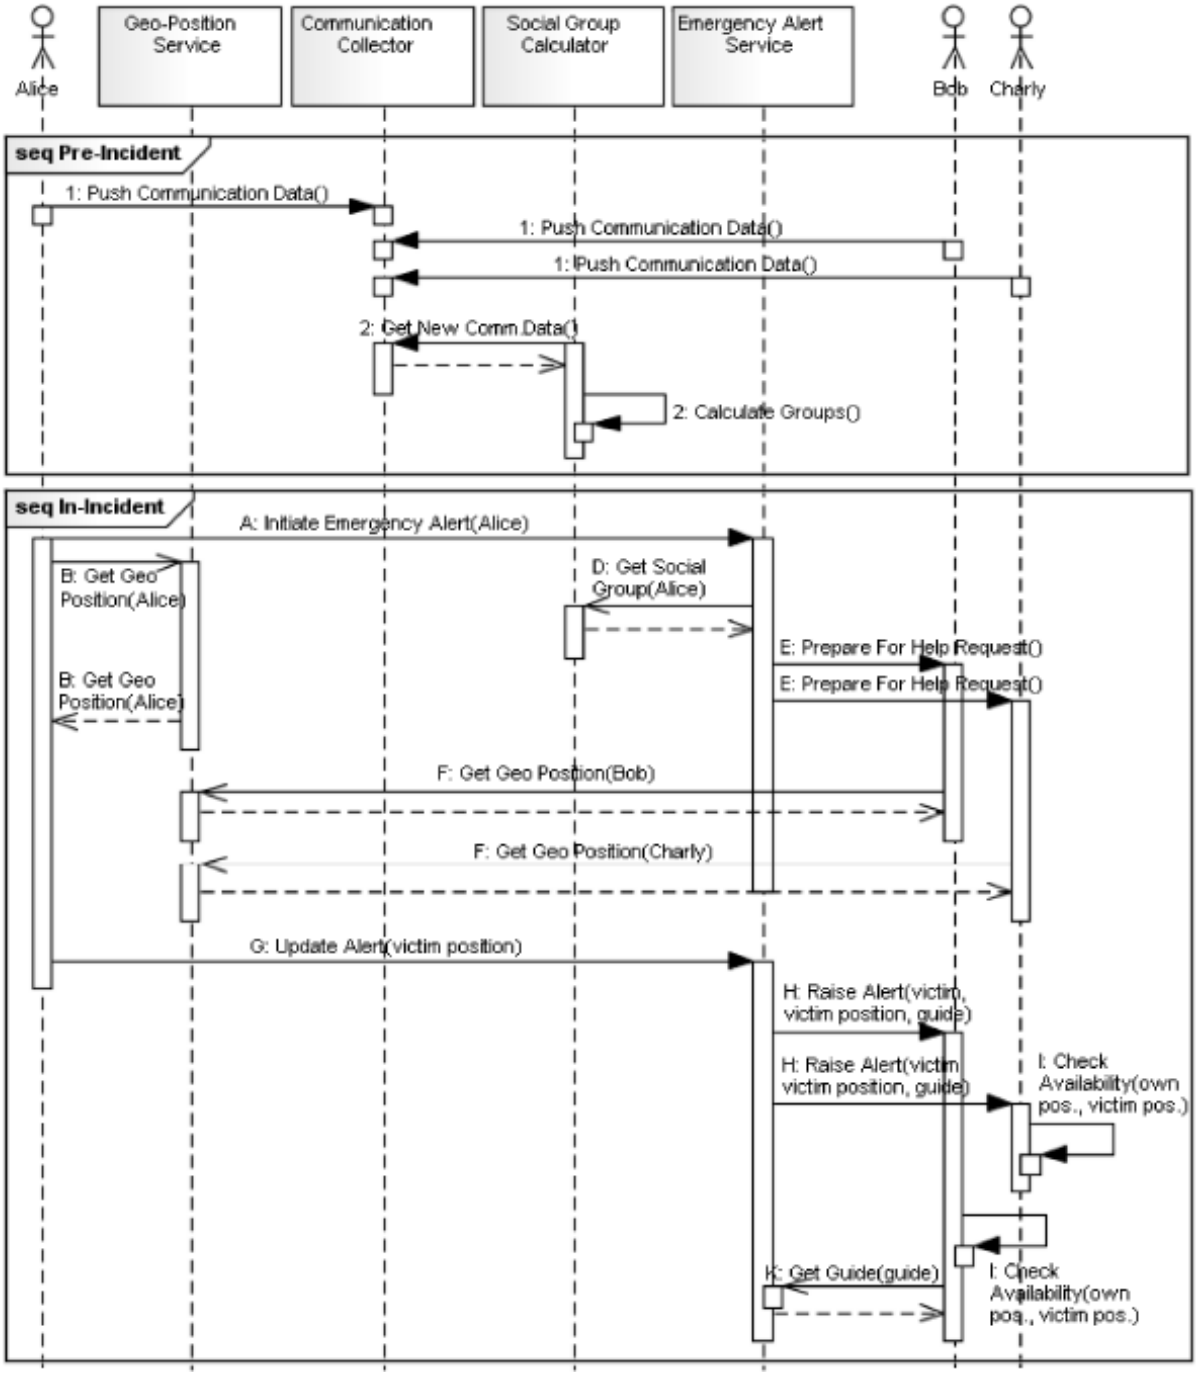

Figure 7: Communication Protocol of a privacy-aware emergency alert service

The privacy-aware architecture is shown in figure 5 (now without geo-position cache). The privacy-aware protocol is designed as follows (see figure 7).

In the **pre-incident** phase only communication data is collected.

The **in-incident** phase starts when a victim presses the alert button. An initiate-alert notice is sent to the Emergency Alert Service (step A) and the victim's phone determines its position (step B). The initiate-alert notice makes the server retrieving the social contacts of the victim (step

D) and sending a prepare-for-alert to all of them (step E). The prepare-for-alert message is a system internal message hidden from the user. This message is sent to decrease the overall time needed to inform a potential helper as the communication is initiated while the victim's GPS unit still determines its geo-position. Now, the phones of all social contacts determine their positions (step F). When the position of the victim has been determined and sent to the Emergency Alert Service (step G), it is broadcasted to the social contact's phones (step H). The phones of the potential helpers will have acquired their own positions by now and check their distance to the victim (step I). When a client has decided that the distance gives a chance to provide help, it retrieves further guidance information from the Emergency Alert Service (step K) and notifies the user.

While knowledge about the social relations of the service participants has still to be processed centrally to identify social groups and to send alert messages to the correct recipients, the server gets only location information from a user who has sent a help request. As the devices of all users do not continuously determine their positions and send location updates to the server, alerts have to be sent to all potential helpers regardless of their current position. Whether a user is close enough to provide help has to be decided by the clients.

1.5.3. Discussion of Design Variants

The centralized system design with the continuous transmission of location updates has two drawbacks. First, continuously determining and transmitting the current position will drain the battery of a smartphone in a few hours. Tests showed that the service drains the battery of our test device in about 10 hours (5 minutes update interval) when the phone is not used otherwise. The service accounted for more than 80% of the energy consumption. Even on more power-efficient devices the battery life will be unacceptable short - especially when considering that the phone will be used, too. A second problem with the centralized approach is that attracting users will be hard because potential participants of the service might not sign-up because of privacy concerns. A geo-position log is able to reveal a lot -- beside spare time activities and favorite shops certainly also personal habits that people want to keep private.

However, a central geo-position cache has also some advantages. A central location database might be desireable for additional security functionalities. For example, a widget that displays the distance to the nearest social contact could convey a sense of security. A geo-position cache also increases the robustness of the system. While the phone signal is often very well even within buildings or narrow streets the GPS signal might not suffice to determine the geo-position. The last-known position (e.g. the position before entering a building) would be helpful in those cases.

1.6. Identifying social groups by clustering social networks

Calling persons willing to help is crucial for the proposed system but their identification is not trivial. Communication networks or 'friend' networks of online social network sites usually contain many links that do not result from close personal relations. Links may connect business partners or co-workers. On social network sites people 'friend' others they rarely know. Therefore, identifying social groups is an approach to separate close personal contacts from other distant contacts that are less willing to help in a case of emergency.

The appropriate cluster algorithms depend on the network that needs to be analyzed. All algorithms need to be highly efficient as the mentioned networks are huge (several million vertices). From communication data weighted networks could be created where an edge connects caller and the callee respectively sender and receiver of a text message. The edges can be weighted by the number of calls or messages. Walk context clustering is a suitable method for this kind of network. It generates overlapping clusters and can reflect that people have several groups of close contacts (family, friends, neighbors) that are almost not connected with each other.

1.6.1. Walk context clustering

Walk context clustering consists of two stages (see figure 9). In the walk stage, a set of restricted random walks is generated by starting a number of walks at each vertex and repeatedly choosing the following vertices randomly among those vertices that are linked by an edge which has a higher weight than the previously taken one (see figure 10).

In the cluster construction stage, clusters get generated from the walks. Walk context clustering assigns a vertex to the cluster of another vertex if both are part of the same walk. A level parameter l specifies the fraction of vertices at the beginning of a walk that are disregarded. The later a pair of vertices appears in a walk the stronger is their connection. The interesting feature of walk-context clustering is that the closeness of two persons can be measured by the maximal level that assigns one person to the cluster of the other one.

```
Data: undirected, weighted graph G = (V, E), constant p
▼ Walk generation
  walkSet ← ∅;
  forall v ∈ V do
    for counter ← 1 to p do
      walk ← (i); last ← 0; i ← v;
      while N = {x|ω_ix > last} ≠ ∅ do
        j ← random element of N;
        last ← ω_ij;
        append j to walk w;
        i ← j;
      walkSet ← walklist ∪ walk;

Data: walkSet ws, vertex v, level l
▼ Cluster Construction
  cluster ← ∅;
  forall w ∈ ws|v ∈ w ∧ pos(v, w) > l do
    forall x ∈ w|pos(x, w) > l do
      cluster ← cluster ∪ x;
```

Figure 9: Algorithm: Walk Context Clustering

A recently developed database-backed update algorithm for the walk stage maintains asymptotically optimal clusters in near real-time (<0.2 sec for a single update on graphs with approximately 500000 nodes and 20 million edges) [41].

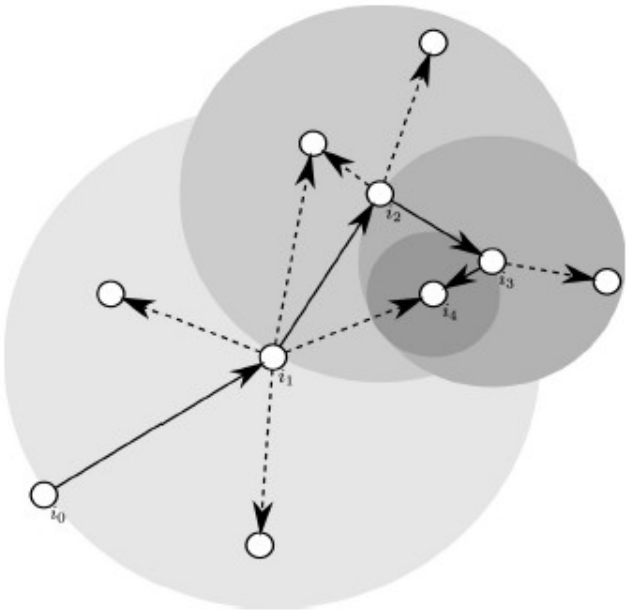

Figure 10: Narrowing search space for successive vertices of the restricted random walk algorithm. The walks terminates when no neighbor is within the search space. Solid arrows symbolize used transitions. Dashed arrows symbolize links to possible successors that have not been chosen by the random process

1.7. Assessment of benefits

1.7.1. Ability to transmit request for help

Emergency situations can result of various incidents, e.g. crime, accidents, medical emergencies. By their nature, accidents happen unexpected and sudden. Some medical emergencies as heart attacks do as well. The ability to make an emergency call in these cases will depend on the physical condition of the person in need.

For victims of violent crime their ability to send an emergency signal depends on the progress of crime. In 25%, respectively 22%, of the robberies analyzed by Smith [42] instant violence or attempts to snatch property don't give the opportunity to call for help. However, in 37% of the robberies the offender approached the victim and demanded money or valuables without immediate violence. In many cases later victims were also aware of an upcoming threat. In those cases it would be possible to send an emergency message.

1.7.2. Chance for nearby help

The helpfulness of the proposed system depends on the availability of close social contacts in the proximity of the site of the emergency. The actual number of persons in ones proximity in general and the number of close contacts with a particular motivation to help will surely depend on several factors, e.g. place and time.

To gain insight into the availability of potential help the MIT reality mining experiment [43] provides an interesting data set. For this experiment a group of 100 persons (75 students or faculty of the MIT Media Laboratory and 25 students of the MIT Sloan Business School) has been equipped with smart phones with special software applications pre-installed. These applications recorded phone numbers of incoming and outgoing calls, text messages, and the id of the cellular tower the phones were connected to during one academic year.

The phone call and text message data from the reality mining data set can be used to construct a communication network. Based on the assumption that the communication intensity of two people reflects the closeness of their relation, social groups can be identified by clustering this weighted network. The availability of nearby help from within the social group of a person in need can be estimated by the number of persons from the same social group whose phones are connected to the same cellular tower as the person in need.

For three consecutive months all communication prior to the specific month has been used to build an undirected, weighted communication network. The edge weights are the number of communication events (calls, text messages). This network has been clustered by the walk context cluster algorithm with the level parameter set to 0.8. The generated clusters had an average size of about 5.

The MIT reality mining data set contains a history of time-stamped connection records of the participating persons. For each month 1000 connection records have been randomly selected. Then, for each persons in a connection record the number of persons of his social group who have been connected to the same cellular tower at the same time have been counted. This simulation showed that on average more than one close fellow student was available for help at any time (see table 1).

	Day (6am-6pm)		Evening (6pm-11pm)		Night (11pm-6am)	
	Group	Others	Group	Others	Group	Others
09/2004	1,16	3,74	1,17	2,73	0,95	1,51
10/2004	1,33	3,12	1,29	3,40	1,14	1,82
11/2004	2,20	3,18	1,53	2,16	1,76	1,03

Table 1: Average number of other persons of same social group connected to same cellular tower for 1000 randomly selected persons and points of time.

E.g. for September 2004 the probability of having at least one person from one's social group in the proximity was 78% - independent of time of the day.

This is just a basic assessment for a particular group of people that has several shortcomings. Due to the lack of more detailed data it was not possible to assess if other social contacts than fellow students were available for help. Proximity could just be estimated by radio cells which have a radius of a few hundred meters in urban areas and a radius up to about 35 km in rural areas. But it is fair to regard the results as an indication that people living in urban areas will usually have at least one of their social contacts in their proximity.

1.7.3. Chance to actually receive help

Once a request for help has been transmitted to a potential helper in the proximity of the site of the emergency the chance to actually receive help depends on the willingness and the ability of the informed persons to get to that place. The ability to get to a specific site can be supported in various ways. For example a map and route directions could be displayed together with the emergency message. The research of Markey [44] showed that also in computer-mediated communication settings help requests that were directed to specific persons had a high probability to receive fast response and the bystander effect was virtually eliminated.

1.8. Conclusions

In this chapter a novel emergency alert service has been introduced which addresses all obstacles to providing help identified in the social emergency intervention process discussed in chapter 1.2.

The service is designed to reduce psychological barriers that result in a bystander effect and inhibit effective help for persons in need. The

analysis of emergency situations and whereabouts of persons in relation to their respective social group suggest that the described service can actually provide a benefit in practice.

Due to its design the service is useful for both scenarios: In the Consumer Social Group scenario the service can be used e.g. in case of a attack on one of travelling friends. It must not be that dramatic: Actually an emergency case may also being lost somewhere in an unknown city or simply not feeling well.

The usage in the Emergency Response scenario is straight forward: The service has been designed for alternative response ways – in detail getting help be the social sphere additionally to public authorities.

However, the following challenges which are beyond the scope of this text are examples of what must be addressed thoroughly for concrete industrial service offerings:

1. Legal issues:
   The service raises e.g. the problem that potential helpers become liable to help and failure to do so may be prosecuted.
2. Privacy:
   The service should be designed in order to minimize the intrusion of privacy of service-subscribers.
3. Emergency dialogue:
   The emergency dialogue could be further automated e.g. by providing an automatic classification of the incident type.
4. Geo-positioning problems: Geo-positioning is still problematic in large buildings, tunnels, subterranean areas (e.g. subway). Enhancements could be based e.g. by embedding geo-position senders in such structures or by image recognition techniques which exploit public geo-coded images of such spaces.

2. Reputation Service

2.1. Introduction

The rapidly increasing popularity of Web 2.0 communities originates in the ease of collaborative content creation and its sharing. More and more community members actively participate in social networks and the community and data growth rates are continuously increasing over time. However, the increasing scale of communities poses the challenge for community platform operators to ensure the quality of content and to prevent violations of laws (e.g. copyright protection, privacy, illegal content) and community rules.

To maintain a desired level of content quality, some platforms have been forced to introduce this form of control, at least temporarily. For example, one could observe with Yahoo! comments and question answering services in Germany in 2009-2010.

Apart from creating undesired content on purpose the core motivations for contributing user generated content to a platform are typically of social nature. Another reason is the human need for sharing knowledge. But it is also motivated by appreciation from other users [45]. Since these motives of users do not necessarily correspond to the purpose of the platform, appropriate user behaviour has to be assured by assessing the value of user contributions. One way to accomplish this is by content moderation [46,47,48]. Since this is a time consuming task it does not scale up with Web 2.0 dimensions. Delays between submission and publishing of new content do not allow for interactive discussions and real time information that is required in emergency response. Beyond this, many community members consider approval of submissions as a form of unwanted censorship. Finally, content moderation is very expensive as people have to be employed to control what is published.

The alternative is to assess the quality of content after it has been published. This requires less administrative resources, but adds uncertainty to the results. In addition, it was already shown that if users know in advance that a quality assessment is performed on their content these users produce content of significantly higher quality [49].

The remainder of this section is organized as follows. This introduction proceeds with a definition of reputation systems, related work and an explanation of the method that is followed for constructing the reputation system. In Section 2.2 this method will be instantiated for the construction of a reputation system for the WKI ER use case. Section 2.3 provides an overview on the architecture of the implementation of the proposed system. Successively conclusion on this Chapter is drawn in Section 2.4.

2.1.1. Reputation Systems

Reputation systems are one way to formalize the assessment of user generated content. In the following, we consider reputation as a quantitative value that explicitly represents a property of an entity of interest. These entities are actors like human users or external agents, transactions, physical objects, or information entities like pictures and texts. Typically these properties of interest are hard to estimate with little effort by examining the entity of interest itself. In the following, we restrict our view to entities that are relevant in Web 2.0 platforms, especially users and user-generated content items. The purpose of a reputation system is to assign a reputation to these entities that has a correlation with a non-visible property like reliability or quality.

The WKI-ER use case implementation allows for collecting and analysing user-generated content. Since this content can be the basis for important decisions made by ER personnel, quality considerations apply especially in this case. The WP4 reputation service developed in this section is an instantiation of a reputation service for improving the quality of user generated content within the WKI ER use case.

2.1.2. Related Work

The process for creating reputation systems is not agreed upon in literature. Comprehensive practical guides like [50] provide valuable insights from practitioners to create reputation systems. They also provide patterns that represent best practices for various aspects of reputations systems, like visualization and usage of scores. There is also evidence that reputation systems have to be carefully designed. Otherwise wrong incentives can be set, so that the overall quality of content decreases, as observed in the SAP Developer Network [51], where the key indicator for reputation has been content quantity instead of content quality.

A broad survey on reputation systems that were described in research papers was performed by Jøsang et.al. [52]. The variety of examined reputation systems illustrates that existing work reputation systems are typically created for specific purposes but often lacks the discussion of purpose and design decisions.

In [53] work is presented that allows ranking resources and users by consideration of user feedback. Although even time dependent decrease of relevance is considered in the model, it lacks a clear definition of purpose and assumptions. Also it does not allow for the integration of additional reputation indicators.

Fujimura and Tanimoto presented in [54] an approach similar to the work presented in this section. They also considered the mutual influence of user reputation and content reputation. Focus in their work lies on algorithmic expression of the reputation system. They do not provide plausible assumptions and indicators for the cold start problem for initial reputation scores.

2.1.3. Method

For developing a suitable reputation system for a user-generated content platform, we follow the method explained in the following.

We assess the (1) Goal of the platform that has to be complemented by the reputation system. These goals can vary from sharing entertaining multimedia content, over recommendations for trips to comments on professional edited news content. After recognizing the goal of the platform, the (2) goal of the reputation system can be derived. Since these goals can also have a wide range like increase of user satisfaction or participation, improvement of content quality, reduction of administration effort. Then a (3) reputation score is defined that quantitatively expresses the targeted goal. Since the reputation score cannot be directly measured, other measurable (4) indications for determining the reputation need to be found, accompanied by (5) assumptions about the the user behaviour. Finally the indicators and the assumptions allow to synthesize a (6) reputation system that is plausibly able to measure the reputation and to support the platform goals.

2.2. Collective Intelligence for Emergency Response

2.2.1. Goal of the WKI-ER system and a supportive Reputation Score

The main usage of the WKI-ER system is the collection of information concerning an ongoing emergency. This information is provided by citizens who are involved or are witnesses, but also from professional personal recruited from various authorities. This information is then annotated, integrated, and further analysed with help of the WKI system. The purpose of these tasks is to provide a full picture of the emergency situation for the responders, to create situation awareness.

Professional responders can be trained to carefully provide the correct information to the right recipient. But the information provided by citizens can be unreliable, because of lack of knowledge or on purpose as vandalism. Therefore the quality of information that is basis for the collective intelligence of the whole WKI-ER system is important. More precisely the important quality aspect of information has to be reliability, because information can influence wide range decisions. Other aspects of quality like relevance, interestingness, novelty, aesthetics, or briefness are either minor considerations or dependent on the current information need of the specific responder. Therefore reliability of information and users will be the core measure that has to be assessed and represented in a reputation system for the WKI ER use case.

2.2.2. Assumptions

For the ER reputation system several assumptions have to be made for the behaviour of users within the system. These are assumptions about the correlation of future behaviour with known properties of users.

1) A user produces content of similar quality and reliability, since his abilities or goals for using the platform do not change. So the known property of reliability of content that was already created by that user can be used for estimating the quality of future content.
2) A user who is well connected with other users through social ties is less likely to create content of low quality, since he has to fear social sanctions by his contacts for misbehaviour.[3] So the known property of connection of the users within the social network of users can be used to estimate the quality of future content.
3) A user who is member in an professional organisation that is known to be reliable will very likely create content of high quality.
4) A user who is reliable is superior in classifying the reliability of specific content than automated categorization. Therefore in case of contradicting information regarding the reliability of a content item, judgement by human users has to be preferred.

2.2.3. Used WeKnowIt Services

The reputation system developed in this chapter integrates different WKI services in order to provide collective intelligence. In the following these services will be briefly introduced.

WP1 Content Rating

The Emergency Response use case implementation allows a user to rate content items that were uploaded to the WKI system. For expressing these ratings along with other information in WP1 there was developed the CURIO ontology[4]. Ratings that are represented in the CURIO vocabulary include properties for the creator of the rating <http://rdfs.org/sioc/ns#has_creator>, the content item <http://rdfs.org/sioc/ns#content>, and the rating score <http://purl.org/net/curio/ns#rate>. This information from the WKI knowledge base can be then represented as rating of reliability for content item c by user u: $r(c,u) \in [0,1]$

WP4 Community Analysis

Social connections between users like explicitly stated friendship or communication can provide an indication for the connectivity of the user

3 This assumption is also supported by the LycosIQ data for quality of user generated content and number of the creators friendships.

4 Burel, Grégoire: CURIO Core Vocabulary Specification v0.4, http://purl.org/net/curio/ns, retrieved 2010-11-30

within the community. WP3 provides several several services[5] for analysing social networks that comprise of social connections between users. The result is a score, like in-degree for every user node. This measure can be then normalized to a value between 0 and 1 measuring the social connectivity for for every user u: $sc(u) \in [0,1]$

WP3 Lexical Spam Detection

Text that can be considered as unwanted and unrelated information typically has properties that can be recognized without deep understanding of its content. The lexical spam detection service of WP3[6] uses statistical learning of n-grams on a corpus from Wikipedia. So unusual patterns in a given text can be classified to potentially have spam characteristics. As result the WP3 Lexical Spam Detection Service returns a value for the spam likeliness for any given textual content c: $sd(c) \in [0,1]$

WP5 Group Management

In an emergency there are normal citizens involved as well as professional organizations that support the response. Examples for professional organizations are fire department, policy department, or public services. The members of these organizations can provide valuable information. The group management service developed in WP5[7] allows for integrating memberships of organizations that are defined outside the WeKnowIt System. When the user is member in an organization that was recognized in advance for being a professional reliable organization. The membership of one user u in one of these organizations can be mapped to a value by the function: $tg(u) \in \{0,1\}$

2.2.4. Reputation System

According to the assumptions and based on the available intelligence an integrated reputation system is developed in the following section.

The core of the WKI-ER reputation system are reputation scores that represent reliability. As shown in Figure 11 the reliability of content items and the reliability of users is measured.

5 WeKnowIt deliverable D.4.1.1 Initial community analysis tools

6 WeKnowIt deliverable D3.1 Prototype of mass question answering

7 WeKnowIt deliverable D5.2.1 Prototypical Knowledge Management Methodology

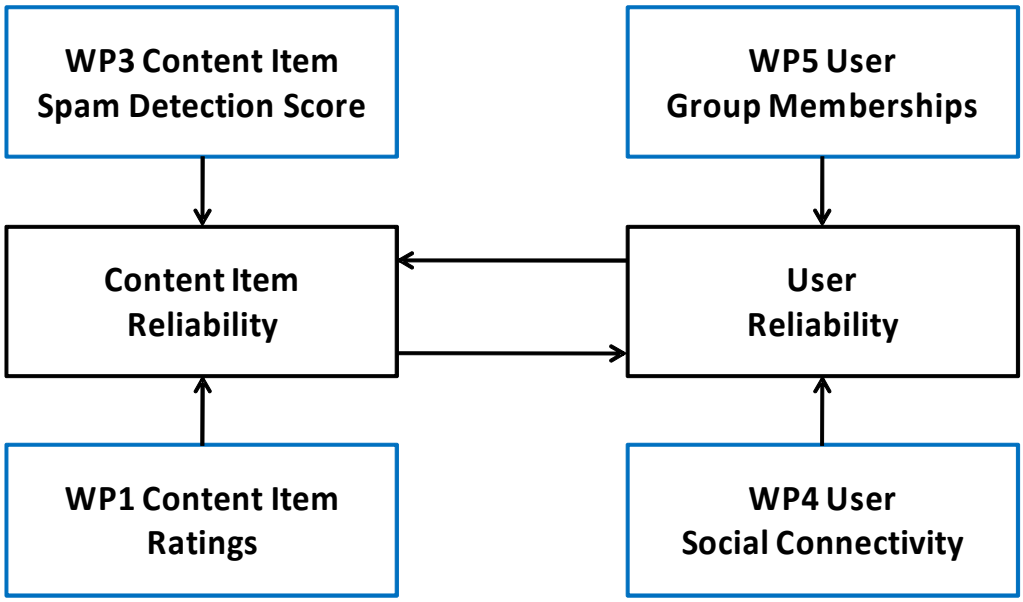

Figure 11: Overview Reputation System

The intelligence that is extracted in other work packages feeds into content item reliability as well as into user reliability. In particular a score that represents the likeliness of a text of being spam (provided by WP3) and therefore unreliable for ER purposes can be used to determine the reliability of this content especially in the absence of other information. If ratings for content items have been provided by other users (collected by WP1) this information will be the basis for calculating the content item reliability. Indication for user reliability is provided by the users social connectivity (determined by WP4) under the assumption that the more connected a user is, the more reliable she might be. A strong indicator for a reliable user is a membership in a group of users that is known in advance to be reliable, like all members of the fire department. This intelligence is provided by WP5. Both content item reliability and user reliability can influence each other. Under the assumption that a reliable user creates reliable content in absence of other information the reliability of the content creator is an indicator of the contents reliability. On the other side when a user creates content that is proven to be reliable this is a strong indication for the the user to be reliable herself.

Basic Elements

The reputation system will measure the reliability for all content items $c \in Content$ that are available in the WKI system and all users $u \in Users$ that are registered in the system.

In the following a set of basic functions will be introduced that allow integrating the intelligence provided by other work packages with a reputation system.

WP1 provides the rating of reliability of content items made by other users. These ratings will be denoted as $r(c,u) \in [0,1]$, mapping a content item c and a user u who has rated c before to the value of the rating, a score between 0 and 1.

WP3 spam detection allows for accessing the the likeliness of a given text of being spam. This spam detection score of a content item c is the value of the function $sd(c) \in [0,1]$, mapping a given content item to a score between 0 and 1, while 1 indicates no spam and 0 indicates spam.

WP4 provides an indication of the social connectivity of a user, based on the number of social ties of this user with other users in the community. This Social connectivity is a mapping of a user u $sc(u) \in [0,1]$ to a score between 0 and 1.

WP5 group management can determine the membership of a user in externally defined groups. When a specific group is in advance considered to be reliable the membership of a user in one reliable group can be expressed as a function $rg(u) \in \{0,1\}$ mapping this user u to a value of 0 or 1.

As additional functions the connection between users and the content the have created will be modelled. In particular the creator of any content item c can be determined by the function $creator(c) \in Users$ while the content item that were created by user u can be determined by $creations(u) \subset Content$.

Default Reliability Scores

The reliability of a content item can be in the best case directly assessed by human users when they review it. But especially new content has no user ratings. Nevertheless the reputation system has to assign a reliability value to the content item. The solution to this is to assign a reliability value based on other known properties. In the case of a content item these are the reliability of the creator of the content item and the spam score if the content item consists of text. Since these indicators for reliability cannot prove reliability directly for the default reliability score of content items the minimum of these both scores is appropriate as default content reliability (dcr):

$$dcr(c) = min(ur(creator(c)), sd(c))$$

A similar approach is taken to determine the user reliability in absence of information about her past contributions. The additional relevant properties that can be considered are the membership in a reliable user group and the social connectivity of the user. The social connectivity is the baseline for that score. If the user is member in a professional user group that is known to be reliable the social connectivity is less relevant, especially since a professional user might not have informal social connections within the platform. Therefore the default user reliability (dur) for a user u is calculated as follows:

$$dur(u) = max(rg(u), sc(u))$$

Calculated Reliability Scores

The reliability score of content items and users is based on the ratings of other users for content.

The calculated content reliability score can be determined as follows if at least one rating for a content item exists. Two cases have to be distinguished. In the case the creator of the content is member of a reliable professional group, there will be always assigned the highest reliability value of 1. In all other cases the reliability of of a content item is determined from the ratings of other users and the reliability of these uses themselves. Since the reliability score a user provides can be considered only so reliable as the user is herself, the calculated content reliability (ccr) is the average of the product of users ratings and their reliability weighted by their reliability:

$$ccr(c)=\begin{cases} 1 & : if\ \ rg(creator(u))=1 \\ \dfrac{\sum_{u \in Users} r(c,u)\, ur(u)^2}{\sum_{u \in Users} ur(u)} & : else \end{cases}$$

Therefore the content reliability (cr) score for any content item c is either the default content reliability (dcr) in cases no ratings exist or the calculated content reliability if there exist ratings for this content item:

$$cr(c)=\begin{cases} dcr(c) & : \text{if no rating for } c \text{ exists} \\ ccr(c) & : \text{else} \end{cases}$$

The calculated reliability score of a user can be determined as follows if at least one rating for a content item created by this user exists. The calculation for this score follows the assumption that the reliability of a users correlates with the reliability of her past content contributions. Therefore in the case the user is not member of a professional group the calculated user reliability (cur) of a user u is the average of of the reliability scores of all content items created by u:

$$cur(u)=\begin{cases} 1 & : if\ \ rg(u)=1 \\ \dfrac{\sum_{c \in creations(u)} cr(c)}{|creations(u)|} & : else \end{cases}$$

If the user is member of a professional group her calculated user reliability is set to 1, since professional users are always considered to be fully reliable.

The final user reliability (ur) of a user u is either the default user reliability (dur) in cases the user has not created an content or the calculated user reliability if there exist content items created by that user:

$$ur(c)=\begin{cases} dur(u) & : \text{if no c} \in creations(u) \text{ exists} \\ cur(u) & : \text{else} \end{cases}$$

Cyclic Dependency of Reliability Scores

The analysis of the formulas for calculating cr and ur shows that each calculation is depends on the other. Consequently for a practical usage and implementation of this reputation system there has to be made a decision how to handle this cyclic dependency. At least two design choices are to be considered in the case when a new rating appears in the system: (1) instant recalculation of all reliability scores, which leads to computationally intensive matrix operations, see e.g. Fujimura et.al., (2) restriction to the recalculation of directly dependant reliability scores.

For an ER application the choice (2) can be preferred. Although the accuracy of the calculated reliability scores is not maximized, the smaller computational effort for calculation is desirable in such time critical applications.

2.3. Architecture and Usage

The architecture of the reputation service is depicted in Figure 12. For each of the WeKnowIt services provided by work packages WP1, WP3, WP4, and WP5 an importer has been developed to convert the inputs and provide the functions for the reputation system. According to the formulas presented in the previous sections, the reputation scores are calculated in the module Reputation Calculator. The calculated scores are then stored with usage of the Reputation Representation module in the provided WP6 Data Storage service.

The reliability scores for users (ur) and content (cr) can be requested by other services through the provided interface of the WP4 Reputation Service.

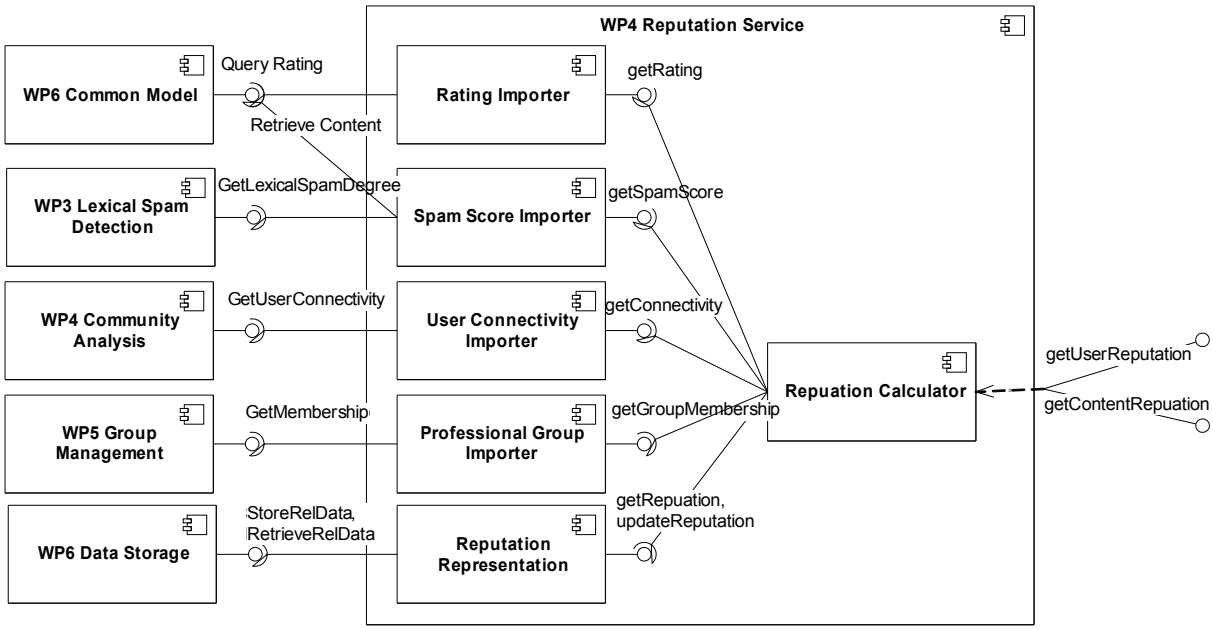

Figure 12: Architecture Reputation Service

The calculated reputation scores are shown to the professional users of the ER system as additional attributes of the existing user-generated content. The reputation of each item of user-generated content can then give the ER personnel additional means of determining the usefulness of the provided information in order to judge the current situation correctly. Reputation as additional information can also become important for decisions of ER personnel in the case of contradicting information.

2.4. Conclusion

In this chapter the need for reputation systems in applications like ER that have a strong need for high quality user generated content was motivated. Based on the application's goals and assumptions about user behaviour a reputation system was derived in order to support ER personnel in assessing the content provided by users of the platform. The reputation system provides reliability scores for users and content items that are derived by usage of intelligence provided by work packages WP1, WP3, WP4, and WP5. Finally the integration of the Reputation Service within the WeKnowIt system was explained.

3. Recommendation service based on Cross-Usage

In this section a service that implements recommendation based on cross-Usage of intelligence is to be presented. As in the previously described services, a key aspect in the proposed service is to enable the user leverage knowledge from the different kinds of intelligence generated in the WKI context so far. To this end, the proposed service aims to exploit the natural multi-modality of social media resources uploaded on the WKI platform by users, analyze and combine intelligence from different modalities and offer recommendations that are based on a joint weighted use of *Media*, *Mass* and *Social Intelligence.* In the end of the section, we show that the proposed service outperforms baseline recommendation techniques that are based on one kind of intelligence alone.

3.1. The proposed framework

Typically, recommender systems have dealt with two dimensions: users and items. In WKI context, items are analogous to resources; however, a number of other dimensions also exist: tags, time, user location. Here, we present a service that uses jointly the aforementioned dimensions and performs user recommendation on the following scenario; a user expresses an interest through his/her tag selections and expects a recommendation of users who share similar interests based on tag assignments, social interactions and photos' content and context. An overview of the service functionality is shown in Figure 13 and a coarse description follows. The users activity in the WKI platform is crawled at specific intervals in an offline process. During that process, the main active topics of interest are extracted based on the uploaded photos and their associated information (i.e. tags, time, etc.). To carry out the topic extraction task, a clustering process is employed, which analyzes a number of dimensions of the uploaded photos (i.e. visual features, tags, social features, etc.) and, in the end, it assigns each cluster to a topic of interest. Each user depending on his/her activity is associated with one or more topics and, thus, topic-based user communities are emerged.

As will be discussed later, the backbone of the service is the clustering process which relies on different kinds of intelligence (i.e. media, mass, social) and combines them in a weighted way, so that: each cluster contains resources that are connected based on one or more kinds of intelligence. For example, one cluster may contain visual-alike images, whereas another cluster may contain images with related tags.

Because the nature of the WKI use-cases encourages sharing and exploration (especially the tourism scenario), user recommender systems are needed to suggest users with similar travelling interests, following same routes, visiting same places, or, even, users who are linked in an

emergency situation. The recommendation service,we propose here, meets the requirements of T4.1, as follows:

- employs multiple kinds of intelligence together, each of which pertains to a different modality of the resources (**exploits the cross-Usage of Intelligence**),
- outputs groups of users that are interested in the same topic as derived by their activity in the WKI platform and their uploaded content (**improves the community knowledge**)

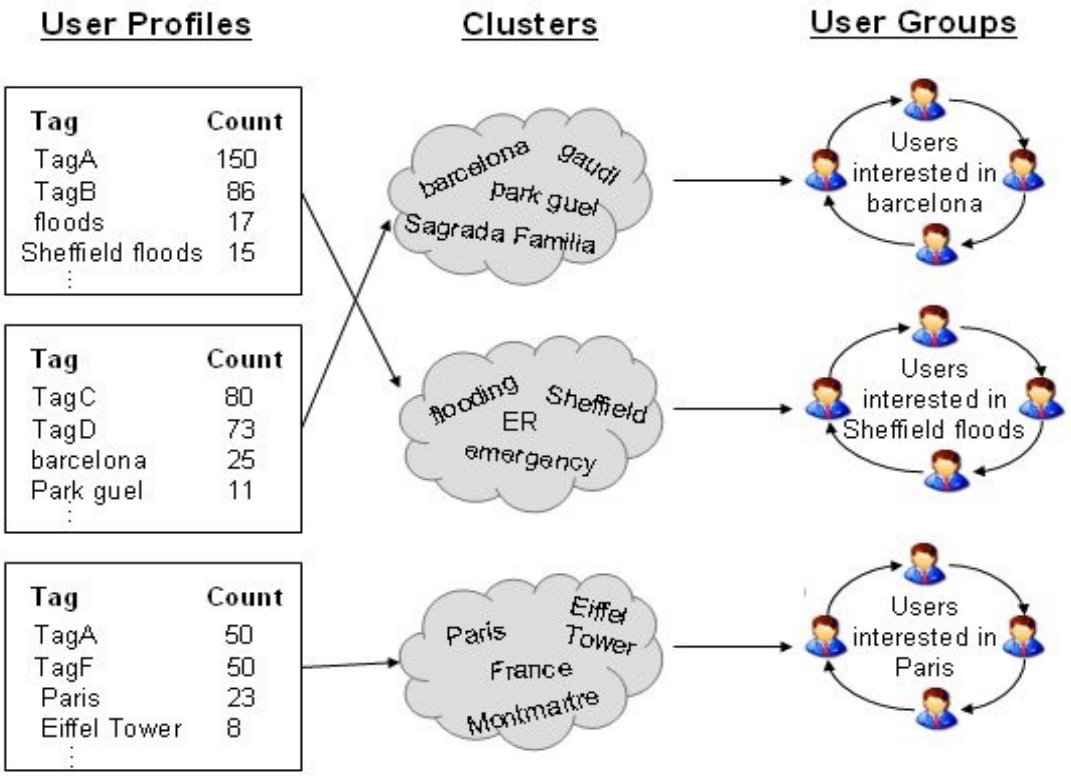

Figure 13: Recommendation service overview

3.1.1. Cross-Usage Motivation

User activity in a social media environment results in correlations of different strength among objects in such environments. For example, a group of visually-alike images have high correlation with respect to their visual features and lower or no correlation in other aspects (*Group A*), whereas images that are visually-alike and are uploaded by a group of connected users show high correlation in terms of both their visual and social connectivity (*Group B*). Understanding the connections among objects and differences associated with the kind of each particular connection may help us interpret more accurately the relationships between users: *Group A* is gathered by users that possibly share the same interest but don't know each other yet, while *Group B* comes from users that are explicitly connected in a social network and, also, share the same

interests. However, such information mining requires methods to measure connectivity (i.e. Similarity) and combine variant types of connectivity together.

In WKI, there have been developed methods for calculating similarities between objects in social media environments by considering various aspects (i.e. modalities) of user contributed and consumed content. Specifically, in WP2 the *Visual-Analysis* service allows to calculate similarities between objects based on the content itself and the *Tag-Processing* service analyzes the tags that accompany the objects and computes object similarities, taking into account tag similarities. A measure of tag similarity is introduced in the WP3 *Local Tag Community Detection* service, which combines information from massive user feedback and extracts patterns of related tags. Moreover, in WP4 *Community Analysis Tool* similarity values between users are extracted, based on existing social networks.

Even though all these kinds of intelligence are not combined naturally with each other, still they carry knowledge about the connections between the resources in WKI platform, with each individual intelligence providing information in a different feature space. However, combination of different kinds of intelligence together is not a straightforward concatenation of individual kinds of intelligence. This is true both in terms of the nature of the raw features in each modality (e.g. sparse, high-dimensional tag co-occurence vectors extracted from tag descriptions, compared to usually dense and low-dimensional descriptors extracted from visual content), as well as in terms of their semantic capacity (e.g. while abstract concepts like "paris" are more easily described with text, concrete objects like "sea" are more easily grounded using visual information). Therefore, techniques are required that will manage to handle the very different characteristics exhibited by the different types of data. Such techniques should aim to learn associations between complex combinations of modality features and semantic concepts. Here, we propose a clustering technique that identifies clusters that exist in different feature subspaces (subspace clustering). The proposed approach analyzes local correlations of data in each feature space (i.e. modality) and assigns weights accordingly. Thus, the extracted clusters lie in different feature subspaces.

As shown in Figure 14, the proposed clustering uses similarity values extracted from WP2 Visual Analysis service, WP2 Tag Processing service, WP3 Tag Community Detection service and WP4 Community Analysis Tool service. The extracted clusters contain resources that exhibit high similarity in weighted combinations of the aforementioned individual similarity values. The proposed clustering will be evaluated on a scenario of user recommendation. Generally, in terms of recommendation it is important to represent as accurately as possible, the connection /correlation between the objects at hand. To this end, the aim of this service in the context of WKI is to demonstrate that the connections

derived from the cross-Usage of intelligence are better than those relying on one kind of intelligence alone.

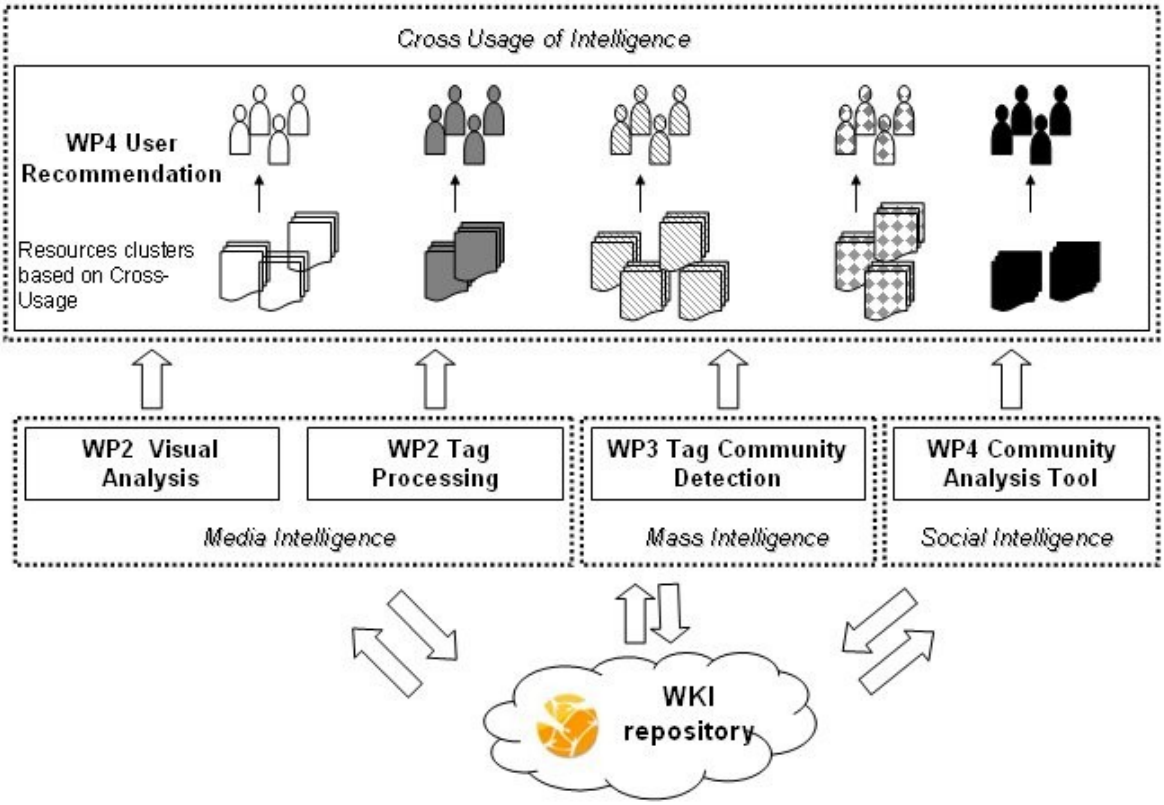

Figure 14: Recommendation service in WKI architecture

3.2. Implementation

As described earlier, there are two main processes in the proposed service: i) the clustering process, and ii) the recommendation process. The processes cooperate in a sequential manner, i.e. user recommendation is based on the outcome of the clustering process.

During the clustering process we are going to tackle the problem of multi-modality in social media, using a feature selection technique. The method we present falls in the category of the so-called subspace clustering approaches that identify clusters of high-dimensional objects and their respective feature subspaces. Specifically, we aim at providing a clustering solution that assigns a feature weight to each dimension on each cluster, based on the correlation of the cluster's objects along the specific dimension. As we will show next, this is a combinatorial optimization problem and we approximate it using the Ant Colony Optimization (ACO) meta-heuristic [55]. Ant colony optimization has been applied successfully to a large number of difficult discrete optimization problems including the travelling salesman problem, the quadratic assignment problem, scheduling, vehicle routing, etc., as well as to routing in telecommunication networks ([56-57]. Although data clustering

techniques for social media have been heavily researched, little research has been dedicated on the use of bio-inspired algorithms for extracting information from this type of content. In this section, we employ an ACO algorithm and demonstrate how it can be used to tackle the problem of combining multi-feature spaces in the clustering of social media problem.

The recommendation process is based on the following scenario: we use the discovered clusters as intermediaries between a user's profile and resources, in order to tailor the results of the clustering process to the user's interests. Thus, users who have uploaded resources that have been assigned in the same cluster are, also, grouped together. So, each user group is perceived as an entity formed through associations and interactions between users in a different feature space. The weighting schemes of each group are considered to establish a recommendation model for users that correlates tags, content, context and social interactions of each user.

First, we are going to describe the ACO metaheuristic and explain how it can be used in a clustering solution. Then, we will give implementation details on the clustering and recommendation algorithms.

3.2.1. ACO for solving a multi-modal clustering

Ant algorithms were inspired by the behavior of real ants when searching for food. When an ant finds food, it releases a chemical substance called pheromone along the path from the food to the nest. Pheromone provides an indirect communication among the ants, since ants typically follow pheromone trails. The amount of pheromone that exists in each path is proportional to the number of ants that have used this path. Pheromone evaporates in time, causing trails and paths that are no longer followed by ants to extinguish.

This pheromone-driven communication between ants have been modeled to solve a number of research problems, one of the most well-known being the Traveling Salesman Problem (TSP). The fact that pheromone evaporates in time causes pheromone trails in longer paths to weaken, as it takes more time for the ants to cross them. On the contrary, a short path is traversed faster, and, thus, the pheromone trail on this path becomes stronger, as it is laid on the path as fast as it can evaporate and many ants follow the trail and reinforce it. This behavior in ant colonies can be used to find shortest paths between nodes in a graph and, thus, providing good solutions to TSP in the following way: Agents are modeled as ants who cross the graph, as they search for food. Initially, the ants are moving randomly. If they meet pheromone trails, they follow them until they visit all the nodes in the graph, constructing, this way, each ant an incremental solution to the problem. When an ant has visited all the nodes, it releases pheromone inversely proportional to the length of the path. Thus, the shorter the path the bigger amount of pheromone is released, attracting other ants to follow the particular path. It is important

to note that pheromone evaporation prevents sticking to local minima and allows a dynamic adaptation when the problem changes.

Now, we are going to show the way the ant colony behavior finding the shortest path can be simulated, in order to solve a clustering problem in multi-feature spaces. The idea is roughly depicted in Figure 15, in which two feature spaces are taken into consideration: i) tag features, and ii) visual features of resources and the description follows:

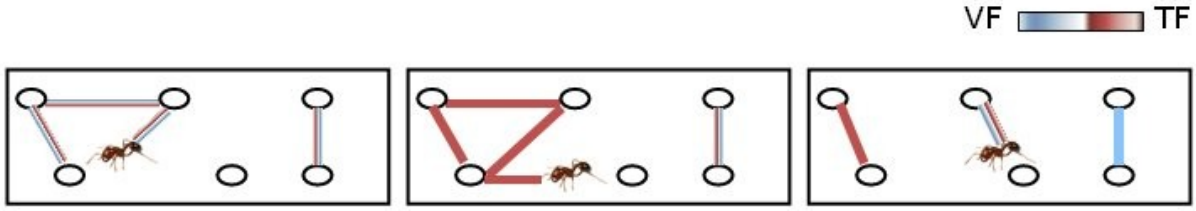

Figure 15: ACO-based clustering using visual and tag features of social media resources: blue paths denote more weight on visual features and red paths denote more weight on tag features

A large number of virtual ants are sent to explore many possible clustering solutions of social media resources. The resources are depicted as graph nodes and the ants join two nodes with an edge if they decide to assign them in the same cluster. The color of the edge shows the weight given by each ant to each feature space individually. The weight is based on the pheromone that there is in each edge. Initially, both feature spaces are given equal weight. Each ant probabilistically assigns each resource to a cluster, based on a measure combining the distance to the cluster center in each feature space individually and the amount of virtual pheromone deposited on the edge to the resource. The ants explore, depositing pheromone on each edge that they cross, until they have all completed their clustering. At this point, the pheromone to each clustering solution is updated (global pheromone updating), so that the edges that have been crossed by many ants become bolder, whereas the remaining ones (that haven't been selected by many ants) become thinner. The amount of pheromone deposited is proportional to the resources correlation along each feature space: the bigger the correlation in tag/visual space, the more pheromone is deposited. The color shows the correlation in each feature space, that is: red color denotes more weight to tag features, whereas blue-colored paths signify clusters that contain objects with high visual similarity.

3.2.2. Resource clustering process

In the context of the ACO-based scenario, described earlier, we present here a resource clustering approach that uses multi-feature spaces and performs subspace clustering, based on the Cross-Usage of Intelligence. As discussed earlier, in WKI each resource is described by a high diversity of features. For instance, an image is associated with the tags that have been assigned to it, the users have uploaded it and seem to like it, the

visual features that describe the visual content of the image, and possibly spatial or temporal information that denote the spatial and temporal context of this particular image. Even though all these facets of information are not combined naturally with each other, still they carry knowledge about the resource, with each facet providing a representation of the particular resource in a different feature space. Thus, to benefit from all this available information, we assume, here, that each resource r is represented by D different feature vectors:

$$r=(F_{1,}F_{2,}\ldots,F_D)$$

where $F_i, 1\leqslant i\leqslant D$ is a feature vector from a corresponding feature space $\mathscr{F}_i$. We can now define the distance between two resources by considering appropriate distance measures for each feature space. For instance, we calculate distances in the tag space, based on tag co-occurrence, whereas we use Euclidean distance to capture the difference in the geographical coordinates between two resources. Thus, given D valid distance measures between the corresponding D feature vectors of the resources r_1 and r_2 we can get their distance as:

$$d^w(r_{1,}r_2)=\sum_{i=1}^{D} w_i\cdot d_i(F_{i1},F_{i2})$$

where d_i is the distance measure employed in feature space $\mathscr{F}_i, 1\leq i\leqslant D$ and w_i is a feature weight that determines the influence of the resources' $i-th$ feature vector to the calculation of the overall distance. In other words, the use of feature weights allows to be given different degree of gravity along each dimension. It holds $\sum_{i=1}^{D} w_i=1$ and $w_i>0.$

The purpose of the proposed approach is to perform clustering on social media resources by optimally combining multi-feature information. They key idea is not to combine all the features together, but to examine local correlations between resources across each dimension and, thus, detect resources' clusters in all feature subspaces. Such techniques are known as subspace clustering and the resulting clusters may refer to different feature subspaces. More formally, we aim at providing solution to the following problem.

Problem Definition: Given a set of N social media resources described by features from D different spaces, a set of D distance measures $d_{1,}d_{2,}\ldots,d_D$ one for each feature space, and an integer $K,$ find a K -partitioning of resources $C_{1,}C_{2,}\ldots,C_K,$ such that $\sum_{i=1}^{K}\sum_{r_{1,}r_2\in C_i} d^w(r_{1,}r_2),$ where d^w a weighted distance in each cluster, is minimized. ■

In order to obtain the $d^w,$ we should define appropriate values for feature weights $w_i, 1\leqslant i\leqslant D,$ so that each feature weight sufficiently captures the local correlation of the resources along the specific dimension. To do so, we employ an ant-inspired algorithm, which is based on ACO metaheuristic. The method of ACO metaheuristic technique which was

proposed by Dorigo is a model of the ant behavior, which is used for combinatorial problems [58]. In the previous section, we presented a general description on how we can modify the ACO algorithm to solve the social media clustering in multiple feature spaces problem. Next, we are going to present a more detailed algorithmic description.

At first, an initialization procedure takes place, during which: i) each ant initializes K centroids randomly. Each centroid c_i is selected to be represented as $(c_{i1}, c_{i2}, \ldots, c_{iD})$, where c_{ij} is a vector from the feature space $\mathscr{F}_j, 1 \leqslant i \leqslant K$, and $1 \leqslant j \leqslant D$ *(lines 7-10)*, ii) the pheromone amount of all graph edges are set to 1 *(line 3)*, iii) the feature weights w_i are set to 0.5, which equals to $1/D$ *(lines 4-6)*, iv) the parameters ρ: pheromone evaporation factor, h: constant used to determine the influence of the distance measure against the pheromone value in the cluster assignment process, are initialized *(line 2)*.

Then, the clustering process begins during which each ant will decide what edges to cross in the graph and what color to paint them. To do so, the following process is repeated for all resources: i) each ant calculates the distance to each cluster centroid in each feature space individually *(lines 15-18)*, ii) considering the feature weights that are already calculated from the previous iteration of the algorithm, each ant estimates an overall distance from each resource to each cluster centroid *(line 19)*, iii) given the overall distance $d^w(r, c)$ (calculated before), the pheromone amount that there is currently on graph edges and the constant h, each ant determines the probability that a resource r should be assigned in a cluster with centroid c, as follows:

$$p(r, c) = \frac{\tau(r, c) \cdot h / d^w(r, c)}{\sum_{i=1}^{K} \frac{\tau(r, c_i) \cdot h}{d^w(r, c_i)}}$$

The ant assigns the resource to the cluster with the highest probability *(line 20)*. This process is illustrated in the graph in Figure 15, as an ant marking an edge between a resource and the other resources already in the cluster. The color of the edge depends on the values of the feature weights w_i.

Having performed the clustering process, new feature weights are calculated for each cluster with centroid c, based on the correlations that there are among the resources in each feature space in the cluster, as follows:

$$w(c, i) = \frac{\sum_{r \in c} d_i(r, c)}{\sum_{r \in c, l=1}^{D} d_l(r, c)}$$

for $1 \leqslant i \leqslant D$ and $d_i(r, c)$ is the distance from the resource r to the cluster centroid c in the feature space $\mathscr{F}_i$ *(line 23-25)*.

Next, new centroids are calculated, based on the assignments in each cluster (line 26). After all ants have done their clustering, the pheromone amount to all solutions is recalculated (lines 30-32). To do so, the quality of each solution needs to be estimated, so that ants that provided good solutions generate more pheromone. The measure we use for ranking the solutions is derived from the definition of clustering, as given in [59] according to which the resources that belong in one cluster should be closely similar to each other, according to some metric of similarity, while the ones that belong to different clusters should be dissimilar. Thus, the most "efficient" ant generates a clustering where: i) in each cluster with centroid c the intra-cluster distance is minimized, that is $IntraDistance_{ant} = min \sum_{r \in c} \sum_{i=1}^{D} w(c,i) \cdot d_{\mathcal{F}_i}$, and ii) the inter-cluster distance is maximized, that is: $InterDistance_{ant} = max \sum_{r_1 \in c_1, r_2 \in c_2, c_1 \neq c_2} \sum_{i=1}^{D} d_{\mathcal{F}_i}$. Thus, we assume that the quality of each solution is given by:

$$Q_{ant} = \frac{InterDistance_{ant}}{IntraDistance_{ant}}$$

Calculating these measures for each solution, we update the current pheromone of each path by considering the total number of ants that have used that path and the quality of their solution. That is:

$$\Delta \tau(r,c) = \sum_{ant} Q_{ant}, \forall ant : (r,c) = 1$$

where $1 \leqslant r \leq N, 1 \leqslant c \leqslant K$. Furthermore, during the global pheromone update, the current pheromone evaporates at a constant rate ρ at each iteration step. The presented ant-based clustering process *(lines 14-32)* is repeated for $NumberOfIteration$ times until the Problem is satisfied.

A pseudocode description of the presented algorithm follows:

```
Algorithm: The ANT-BASED clustering algorithm
combining features from multi-feature spaces
1:  W_F(0) ← Generate-Initial-Weighting
2:  C(0) ← Generate-Initial-Clustering ( W_F(0) )
3: for t←1 to NumberOfIteration do:
4:   Evaluation();
5:   Update-Pheromone( W_F(t) , C(t) );
6:   if stopping−criteria−met then exit for;
7:     C(t+1)← Generate-Clustering( W_F(t) )
8: end;
9: return C(t)
```

3.2.3. User Recommendation Process

The recommendation process is based on the results of the ACO-based clustering. More specifically, the discovered clusters act as intermediaries between a user's profile and resources, in order to tailor the results of the clustering process to the user's interests. This is done as follows: Each resource cluster is used to specify relationships among users. Since, typically, a resource group is an association of dependent or related resources, the users that have uploaded resources of the same cluster are expected to share similarities in their profiles as well. Thus, they are grouped together in the same cluster and these user clusters serve as the basic units in making recommendations.

```
Algorithm: The recommendation algorithm based on ANT-
BASED clustering for a user u_0
1:  C^R_F  ← Generate-ANT-BASED-Clustering ( R,F )
2:  CT  ← Generate-Cluster-Topics ( C^R_F )
3:  C^U_F  ← Generate-User-Clustering
4:  search−topic  ←  Initialize-Search-Topic;
5: for user u in the same-cluster with u_0 do:
6:   Sim( u,u_0 ) ← Calculate-Similarity ( search−topic );
7: end;
9: return  k nearest users;
```

3.3. Experimental Results

In order to test the described ACO clustering algorithm, a dataset from Flickr online photo management and sharing application was crawled. As we were interested, initially, to check the functionality of the algorithm, we conducted experiments on a rather small dataset and examine how the ant-based clustering algorithm performs better than a typical clustering algorithm that relies on predefined feature spaces. Thus, the dataset was restricted to 3000 images (size 500x735) that depict cityscape, seaside, mountain, roadside, landscape and sport-side locations[8].

At first, we applied typical clustering algorithms (K-Means, Hierarchical and Cobweb) for various values of *K*, based only on tag features. The distance measure used is a combination of tag co-occurence and WordNet distance [60]. We examine the extracted clusters manually and observed that many clusters had poor accuracy, i.e. they contained resources not depicted related themes, although sharing related tags. Figure 16 shows

8 For Flickr resources and metadata download the Flickr API along with the utility wget were used

some indicative snapshots of this type of clustering. Specifically, the limitation of algorithms relying solely on tag information to handle ambiguous terms is shown. A snapshot of a cluster containing social media resources about *Paris* is shown in (b).

(a) **(b)**

Figure 16: Indicative outcome of clustering using tag features (a) Cluster around the ambiguous tag *rock,* (b) Cluster around the tag *Paris*

To minimize the intrinsic shortcomings of tagging features, we embedded in the clustering process the visual features of the images. More specifically, we performed clustering using both tag and visual features in a sequential way: first we apply tag-based clustering and then clustering based on the visual features of the resources [71]. The visual feature extraction was based on the MPEG-7 standard [61], which defines appropriate descriptors together with their extraction techniques and similarity matching distances.

More specifically, the MPEG-7 eXperimentation Model, XM provides a reference implementation which was utilized in our approach [62]. The descriptors used were the Color Structure Histogram (*CSH)* and Edge Histogram *(EH)* descriptors, chosen due to their effectiveness in similarity retrieval. Their extraction was performed according to the guidelines provided by the MPEG-7 XM and then, an image feature vector was produced, for every resource, by encompassing the extracted MPEG-7 descriptors in a single vector. Thus, typical clustering algorithms could be applied, using the distance functions that are defined in MPEG-7 XM [71]. A set of users checked the extracted clusters manually and assessed their quality. The results showed that in many cases the accuracy was improved. This was especially true for clusters that contained ambiguous terms (e.g. rock - stone, rock -music). Especially for cases that the two senses of the ambiguous tag differed a lot visually, the algorithm succeeded to distinguish the different senses of the ambiguous tag, by dividing the corresponding resources into different clusters (cf. Figure 17).

However, there were cases that the combination of tag and visual features worsen the clustering outcome. An example of such case is the *Paris* cluster and, generally, clusters whose topic is not related to a particular visual representation (abstract concepts). Indeed, we saw that a cluster describing *Paris* was extracted based on tag features of resources. This cluster no longer exists, if we employ the visual features to the clustering. As shown in Figure 17, this type of clustering succeeds in assigning resources with uniform visual appearance together. This way though it misses clusters of abstract resources that refer to the same topic but differ visually.

The aforementioned example shows that the equal consideration of diverse features describing a resource does not always yield the optimal results. Some cluster resources can be only extracted by using a specified combination of features. We apply the presented ACO-based method using two modalities of the images, i.e. textual and visual. For the representation of images in the tag space we used the approach described in [60] and for their representation in a visual space we used their Color Structure Histogram *(CSH)* and Edge Histogram *(EH)* descriptors, as described above. We conducted a number of experiments with different values of the parameters ρ and h. Figure 18 shows indicative clustering results of the ant-inspired algorithm. It can be seen that the algorithm sufficiently captures clusters in different feature subspaces and it managed to handle the *Paris* cluster well.

(a) **(b)**

Figure 17: Indicative outcome of clustering using tag features and visual features (a) Cluster around the topic rock (sense:stone), (b) Cluster around the topic rock (sense:music)

(a) **(b)**

Figure 18: Indicative outcome of clustering using presented ACO algorithm (a) Cluster around the topic *sea,* $w_{TF}=0.136,\quad w_{VF}=0.86,\quad \rho=0.85,\quad h=0.75,$ **(b) Cluster around the topic *Paris*,** $w_{TF}=0.84,\quad w_{VF}=0.002,\quad \rho=0.85,\quad h=0.75,$

The presented experiments are preliminary and were performed to evaluate qualitatively the use of ACO methods in social media clustering. We reached useful conclusion that sum up as follows. It is apparent that the restriction in one feature space deprives information that can come out handy in the task of social media clustering. On the other hand, by considering all the feature spaces (i.e. modalities in social media) equally important we may miss clusters of related object that have similarity using a specific combination of features.

We assume that users that have uploaded resources of the same cluster share similarities in their profiles as well. Thus, having obtained clusters of social media, corresponding user clusters are formed with users having uploaded resources in each cluster. An interesting feature of the user clusters is that the connections in each cluster are weighted and the weights are the ones that were extracted for the corresponding social media cluster. For instance, some users may be connected because they use similar tags, whereas others may be connected because they use similar tags and have also strong relationship through their social networks. Such user clusters serve as the basic units in making recommendations in WKI context, and, in addition, we claim that the proposed analysis can be used to:

- Enrich the social graph of communities by utilizing multiple kinds of intelligence
- Improve the recommendation process compared to recommendations based on one kind of intelligence alone

In this section we proposed applying subspace clustering for obtaining clusters of social media in various feature subspaces. An ant-inspired

algorithm was presented to realize this task. The outcome of this algorithm was utilized for performing user recommendations. Future work involves the more controlled testing and evaluation of the recommendations on specific users who will rate the recommendation outcome, as well as the testing of the algorithm on larger datasets.

3.4. Related Work

3.4.1. Mining user similarity

Mining similarity information may be useful for numerous scenarios. It may be used for information discovery, by making users aware of people who share similar interests and who may be commenting on interesting blogs or bookmarking interesting articles. It may be used in expertise location scenarios where an expert is not available, but people with similar expertise may be approached [63]. Recommender systems already make use of similarity (in collaborative filtering [64]), and may gain from expanding their similarity information sources beyond their own system, and from better understanding the characteristics of different similarity sources. Promoting response for advice is another motivation for identifying and highlighting similar people: Constant et al. In [65] discuss the "kindness of strangers" and argue that people are likely to provide help to people who are similar to them. Homophily, a term coined by Lazarsfeld and Merton in [66], refers to the tendency of people to associate and bond with others who are similar to them. Other scenarios for leveraging user similarity information include choosing group members [67], building and maintaining a community [68], and clustering of similar users to better understand their behaviors in social applications [69].

3.4.2. Social media applications with multi-modal analysis

The multi-modal analysis has been used as the core analysis component of various different applications in social media. First of all, the fact that most people are fond of uploading and sharing content in social media environments motivates research efforts towards better browsing and retrieval functionalities in such environments. Furthermore, the intrinsic limitations of these systems (e.g. tag ambiguities, erroneous metadata or lack of metadata, etc.) addresses the need for exploitation of features in multi modalities, in order to provide users with as much possible relevant content. To this end, in [70], the authors present a method for a social image database browsing and retrieval by exploiting both tag and visual features of the images in a supplementary way. Indeed, it is shown that the visual features can support the suggestion of new tags and contribute to the emergence of interesting (semantic) relationships between data sources. Through the use of a navigation map, these emergent relationships between users, tags and data may be explored.

Another approach in this direction is met at [71], where the authors present clustering algorithms that improve the retrieval in social media by exploiting features from multi spaces. Specifically, a two-step clustering approach is proposed that uses tag features at the 1st step, in order to get resources clusters that refer to certain topics, and visual features at the 2nd step, in order to further "clean" and improve the precision to the extracted clusters. The use of both visual and text features is also described in [72], where the authors deploy the visual annotations, also known as "notes" in Flickr, and it is shown that the retrieval of social media content improves significantly by combining tags and visual analysis techniques.

A number of works have addressed the problem of identifying photos from social media environments that depict a certain object,location or event. In [73] they analyze location and temporal features from geotagged photos from Flickr, in order to track tags that have place semantics (i.e. they refer to an object in a restricted location) or event semantics (i.e. they are met in specified time periods). Then, they employ tag-based clustering on these specific tags, followed by clustering on their visual features, in order to capture distinct viewpoints of the object of interest. The same authors in [74] combine tags with content-based features and analysis techniques, in order to get groups of music events photos. Likewise, in [75-76] the authors use various modalities of photos (i.e. visual, textual, spatial, temporal proximity), in order to get photo collections in an unsupervised fashion. Apart from the obvious retrieval application, the outcome of the described methods that perform object or POI identification can be used for training of multimedia algorithms, whereas these methods that extract social media content associated with particular events can be exploited for faceted browsing of events and related activities in browsers.

Most of the aforementioned methodologies can be exploited for tag recommendations in the sense that they extract tags associated to a particular event, object, location or, in general, cluster of related resources. The problem of tag recommendation has been further studied in [77], where the authors suggest an approach for recommending tags by analyzing existent tags, visual context and user context in a multimedia social tagging system. Tag recommendation techniques were, also, proposed in [78], where the authors suggest four methods for ranking candidate tags and in addition, they present the semantics of tags in Flickr.

4. Conclusion

In this document we describe three services developed with Work package 4 Social Intelligence.

The **Emergency Alert Service** has been designed and developed to break the bystander effect and alert contacts of a victim's social sphere in parallel to public authorities. It is a location-based service that activates nearby members of the social group of a victim in the case of an emergency. It is a mobile service which can be installed as an applet on smart phones (e.g. Android based phones). When activated it works as an emergency call agent and informs social contacts (friends, family, and colleagues) and public authorities about the emergency situation. Current geo-position and routing information is provided. The service is designed privacy-aware and dynamic. Due to its design it can be used in both scenarios, the emergency response scenario as well as the consumer social group study.

The second service is the **Recommendation Service** (RecS) which has been designed to reveal connections between users, based on their activities in the WKI platform. Its backbone is a clustering process which identifies clusters of resources in different feature spaces. The main difference to the typical recommendation algorithms is that it differentiates between users' interconnections and assigns weights on each kind of similarity they share i.e. social similarity (through their social networks), content similarity (through tags and photos). To estimate the weights, the service makes usage of intelligence provided by WP2, WP3 and WP4.

The third service described is the **Reputation Service** (RepS). It is a service that supports the generation of high quality user generated content in the ER use case. Based on the application's goals and assumptions about user behaviour this reputation system was derived in order to support ER personnel in assessing the content provided by users of the platform. The reputation system provides reliability scores for users and content items that are derived by usage of intelligence provided by work packages WP1, WP3, WP4, and WP5.

[1] Spiegel Online: Beaten to death in broad daylight - germany shocked by brutal commuter train murder (2009), http://www.spiegel.de/international/germany/0, 1518,649134,00.html

[2] Spiegel Online: Tödlicher Angriff: Mehrere Menschen beobachteten Münchner S-Bahn-Attacke (Sep 2009), http://www.spiegel.de/panorama/justiz/0,1518, 649288,00.html

[3] Spiegel Online: Notruf-Mitschnitt belegt Brutalität der S-Bahn-Prügler (Sep 2009), http://www.spiegel.de/panorama/gesellschaft/0,1518,656858,00. html

[4] Latane, B., Darley, J.: The Unresponsive Bystander: Why doesn't he help? Appleton-Century-Crofts, NY (1970)

[5] Yuan, Y., Detlor, B.: Intelligent mobile crisis response systems. Communications of the ACM 48(2), 95 – 98 (2005)

[6] Faraj, S., Xiao, Y.: Coordination in fast-response organizations. Management Science 52(8), 1155 – 1169 (8 2006)

[7] Chen, R., Sharman, R., Rao, H.R., Upadhyaya, S.J.: Coordination in emergency response management. Communications of the ACM 51(5), 66 – 73 (2008)

[8] Comfort, L.K., Kilkon, K., Zagorecki, A.: Coordination in rapidly evolving disaster response systems: The role of information. American Behavioral Scientist 48(3), 295 – 313 (2009)

[9] Palen, L., Liu, S.B.: Citizen communication in crisis: Anticipating a future of ict- supported public participation. In: CHI 2007 Proceedings. pp. 727 – 736. ACM (2007)

[10] Darley, J.M., Latane, B.: Bystander intervention in emergencies: Diffusion of responsibility. Journal of Personality and Social Psychology 8(4), 377 – 383 (1968)

[11] Milgram, S.: The experience of living in cities. Science 167, 1461 – 1468 (1970)

[12] Clark, R., Word, L.E.: Why don't bystanders help? because of ambiguity? Journal of Personality and Social Psychology 24, 392 – 400 (1972)

[13] Piliavin, I.M., Piliavin, J.A., Rodin, J.: Costs, diffusion, and the stigmatized victim. Journal of Personality and Social Psychology 32(3), 429 – 438 (1975)

[14] Shotland, R.L., Straw, M.K.: Bystander response to an assault: When a man attacks a woman. Journal of Personality and Social Psychology 34, 990 – 999 (1976)

[15] Miller, D.T., McFarland, C.: Pluralistic ignorance: When similarity is interpreted as dissimiliarity. Journal of Personality and Social Psychology 53, 298 – 305 (1987)

[16] Monin, B., Norton, M.: Perceptions of a fluid consensus: Uniqueness bias, false consensus, false polarization, and pluralistic ignorance in a water conservation crisis. Personality and Social Psychology Bulletin 29, 559 – 567 (2003)

[17] Garcia, S.M., Weaver, K., Moskowitz, G.B., Darley, J.M.: Crowded minds: The implicit bystander effect. Journal of Personality and Social Psychology 83, 843 – 853 (2002)

[18] Rutkowski, G.K., Gruder, C.L., Romer, D.: Group cohesiveness, social norms, and bystander intervention. Journal of Personality and Social Psychology 44, 545 – 552 (1983)

[19] Baumeister, R.F., Chesner, S.P., Sanders, P.S., Tice, D.M.: Who's in charge here? group leaders do lend help in emergencies. Personality and Social Psychology Bulletin 14, 17 – 22 (1995)

[20] Shotland, R.L., Heinold, W.D.: Bystander response to arterial bleeding: Helping skills, the decision-making process, and differentiating the helping response. Jour-nal of Personality and Social Psychology 49, 347 – 356 (1985)

[21] Cramer, R.E., McMaster, M.R., Bartell, P., Dragna, M.: Subject competence and the minimization of the bystander effect. Journal of Applied Social Psychology 18, 1133 – 1148 (1988)

[22] Schwartz, S.H., Gottlieb, A.: Bystander anonymity and reaction to emergencies. Journal of Personality and Social Psychology 39, 418 – 430 (1980)

[23] Darley, J.M., Batson, C.: "From Jerusalem to Jericho": A study of situational and dispositional variables in helping behavior. Journal of Personality and Social Psychology

27(1), 100 – 108 (1973)

[24] Batson, C.D., Cochran, P.J., Biederman, M.F., Blosser, J.L., Ryan, M.J., Vogt, B.: Failure to help when in a hurry: Callousness or conflict? Personality and Social Psychology Bulletin 4, 97 – 101 (1978)

[25] Stürmer, S., Snyder, M., Omoto, A.M.: Prosocial emotions and helping: The moderating role of group membership. Journal of Personality and Social Psychology 88(3), 532 – 546 (2005)

[26] Levine, M., Prosser, A., Evans, D., Reicher, S.: Identity and emergency intervention: How social group membership and inclusiveness of group boundaries shape helping behavior. Personality and Social Psychology Bulletin 31(4), 443 – 453 (2005)

[27] Nadler, A., Halabi, S.: Intergroup helping as status relations: Effects of status stability, identification, and type of help on receptivity to high-status group's help. J. of Personality and Social Psychology 91(1), 97 – 110 (2006)

[28] Moriarty, T.: Crime, commitment, and the responsive bystander: Two field experiments. Journal of Personality and Social Psychology 31, 370 – 376 (1975)

[29] Shotland, R.L., Stebbins, C.A.: Bystander response to rape: Can a victim attract help? Journal of Applied Social Psychology 10, 510 – 527 (1980)

[30] Markey, P.M.: Bystander intervention in computer-mediated communication. Computers in Human Behavior 16, 183 – 188 (2000)

[31] Hoser, B., Geyer-Schulz, A.: Eigenspectral analysis of hermitian adjacency matrices for the analysis of group substructures. The Journal of Mathematical Sociology 29(4), 265 – 294 (2005)

[32] Granovetter, M.: The strength of weak ties: a network theory revisited. In: Marsden, P.V., Lin, N. (eds.) Social Structure and Network Analysis, chap. 5, pp. 105 –130. Sage Publications, Inc., Beverly Hills, California (1982)

[33] Granovetter, M.S.: The strength of weak ties. The American Journal of Sociology 78(6), 1360 – 1380 (1973)

[34] Rungg, A.: Privatspäre 2.0. Financial Times Deutschland 17.2.2010, p. 25 (2010)

[35] Kikiras, P., Drakoulis, D.: An integrated approach for the estimation of mobile subscriber geolocation. Wireless Personal Communications 30(2), 217–231 (2004)

[36] Borkowski, J., Lempiainen, J.: Practical network-based techniques for mobile positioning in UMTS. EURASIP J. Appl. Signal Process. 2006, 149–149 (2006)

[37] EU: Directive 2006/24/EC of the European Parliament and of the Council on the retention of data generated or processed in connection with the provision of publicly available electronic communications services or of public communications networks and amending Directive 2002/58/EC (Mar 2006)

[38] Ovelgönne, M., Geyer-Schulz, A.: Cluster cores and modularity maximization. In: ICDMW '10. IEEE International Conference on Data Mining Workshops (2010), to appear.

[39] Ovelgönne, M., Geyer-Schulz, A., Stein, M.: Randomized greedy modularity optimization for group detection in huge social networks. In: SNA-KDD'10: Proceedings of the 4th Workshop on Social Network Mining and Analysis. ACM, New York, NY, USA (2010), to appear.

[40] Ovelgönne, M., Sonnenbichler, A.C., Geyer-Schulz, A.: Social emergency alert service - a location-based privacy-aware personal safety service. In: NGMAST '10. Proceedings of the 2010 Fourth International Conference on Next Generation Mobile Applications, Services and Technologies. pp. 84–89. IEEE Computer Society (2010)

[41] Franke, M., Geyer-Schulz, A.: An update algorithm for restricted random walk clustering for dynamic data sets. Advances in Data Analysis and Classification 3(1), 63 – 92 (2009)

[42] Smith, J.: The nature of personal robbery. Home Office Research Study 254, Home Office Research, Development and Statictics Directorate (2003)

[43] Eagle, N., Pentland, A.: Reality mining: Sensing complex social systems. Personal and Ubiquitous Computing 9, 1 – 14 (2005)

[44] Markey, P.M.: Bystander intervention in computer-mediated communication. Computers in Human Behavior 16, 183 – 188 (2000)

[45] Daugherty, T.; Eastin, M. & Bright, L. Exploring consumer motivations for creating user-generated content Journal of Interactive Advertising, 2008, 8, 1-24

[46] Lampe, C. & Resnick, P. Slash (dot) and burn: distributed moderation in a large online conversation space Proceedings of the SIGCHI conference on Human factors in computing systems, 2004, 543-550

[47] Heymann, P.; Koutrika, G. & Garcia-Molina, H. Fighting spam on social web sites: A survey of approaches and future challenges Internet Computing, IEEE, IEEE, 2007, 11, 36-45

[48] Schwagereit, F.; Sizov, S. & Staab, S. Finding Optimal Policies for Online Communities with CoSiMo Proceedings of the WebSci10: Extending the Frontiers of Society On-Line, April 26-27th, 2010, Raleigh, NC: US, 2010

[49] Cosley, D.; Frankowski, D.; Kiesler, S.; Terveen, L. & Riedl, J. How oversight improves member-maintained communities Proceedings of the SIGCHI conference on Human factors in computing systems, ACM, 2005, 11-20

[50] Farmer, F. & Glass, B. Treseler, M. E. (Ed.) Building web reputation systems Oreilly & Associates Inc, 2010

[51] Kwan, M. & Ramachandran, D. Golbeck, J. (Ed.) Trust and Online Reputation Systems Computing with Social Trust, Springer London, 2009, 287-311

[52] Jøsang, A.; Ismail, R. & Boyd, C. A survey of trust and reputation systems for online service provision Decision Support Systems, Elsevier, 2007, 43, 618-644

[53] Gulli, A.; Cataudella, S. & Foschini, L. Tc-socialrank: Ranking the social web Algorithms and Models for the Web-Graph, Springer, 2009, 5427/2009, 143-154

[54] Fujimura, K. & Tanimoto, N. The EigenRumor algorithm for calculating contributions in cyberspace communities Trusting Agents for Trusting Electronic Societies, Springer, 2005, 3577/2005, 59-74

[55] Dorigo, M.: Optimization, Learning and Natural Algorithms. Ph.D. thesis, Politecnico di Milano, Italy (1992)

[56] Caro, G.D., Ducatelle, F., Gambardella, L.M.: Anthocnet: an adaptive nature-inspired algorithmfor routing in mobile ad hoc networks. European Transactions on Telecommunications 16(5), 443–455 (2005)

[57] Blum, C.: Ant colony optimization: Introduction and recent trends. Physics of Life Reviews 2, 353–373 (2005)

[58] Dorigo, M., Caro, G.D.: The ant colony optimization meta-heuristic (1999)

[59] Xu, R., Wunsch, I.: Survey of clustering algorithms. Neural Networks, IEEE Transactions on 16(3), 645–678 (2005)

[60] E. Giannakidou, V. Koutsonikola, A. Vakali and I. Kompatsiaris, "Co-clustering Tags and Social Data Sources", In Proc. 9th International Conference On Web-Age Information Management (WAIM' 2008), IEEE Computer Society, pp 317-324, July 20-22, 2008, Zhangjiajie, China

[61] Manjunath, B.S., Ohm, J.R., Vinod, V.V., Yamada, A.: Colour and texture descriptors. IEEE Trans. Circuits and Systems for Video Technology, Special Issue on MPEG-7 11(6), 703–715 (2001)

[62] MPEG-7 Visual Experimentation Model (XM). Version 10.0, ISO/IEC/JTC1/SC29/WG11, Doc. N4062 (2001)

[63] Balog, K. & de Rijke, M. 2007. Finding similar experts. Proc. SIGIR '07, 821-822

[64] Goldberg, D., Nichols, D., Oki, B. M., & Terry, D. Using collaborative filtering to weave an information tapestry. Commun. ACM 35, 12 (Dec. 1992), 61-70

[65] Constant, D., Sproull, L., & Kiesler, S. 1996. The kindness of strangers: the usefulness of electronic weak ties for technical advice. Organization Science 7 (2), 119–135

[66] Lazarsfeld, P. F. & Merton, R. K. Friendship as a social process: A substantive and methodological analysis. Freedom and Control in Modern Society (1954), 18-66.

[67] Hinds, P. J., Carley, K. M., Krackhardt, D., & Wholey, D. 2000. Choosing work group members: Balancing similarity, competence, and familiarity. OBHDP 81, 2 (2000), 226–251

[68] Sumi, Y. & Mase, K. 2000. Supporting awareness of shared interests and experiences in community. Proc. Group'00, 35-42

[69] Xiao, J., Zhang, Y., Jia, X., & Li, T. 2001. Measuring similarity of interests for clustering web-users. Proc. ADC'01, 107-114

[70] Aurnhammer, M., Hanappe, P., Steels, L.: Augmenting navigation for collaborative tagging with emergent semantics. In: International Semantic Web Conference (2006)

[71] Giannakidou, E., Kompatsiaris, I., Vakali, A.: Semsoc: Semantic, social and content-based clustering in multimedia collaborative tagging systems. In: ICSC, pp. 128–135 (2008)

[72] Olivares, X., Ciaramita, M., van Zwol, R.: Boosting image retrieval through aggregating search results based on visual annotations. In: Proceeding of the 16th ACM international conference on Multimedia, MM '08, pp. 189–198. ACM, New York, NY, USA (2008)

[73] Kennedy, L.S., Naaman, M., Ahern, S., Nair, R., Rattenbury, T.: How flickr helps us make sense of the world: context and content in community-contributed media collections. In: ACM Multimedia, pp. 631–640 (2007)

[74] Kennedy, L., Naaman, M.: Less talk, more rock: automated organization of community contributed collections of concert videos. In: Proceedings of the 18th international conference on World wide web, WWW '09, pp. 311–320. ACM, New York, NY, USA (2009)

[75] Quack, T., Leibe, B., Gool, L.J.V.: World-scale mining of objects and events from community photo collections. In: CIVR, pp. 47–56 (2008)

[76] Crandall, D.J., Backstrom, L., Huttenlocher, D., Kleinberg, J.: Mapping the world's photos. In: Proceedings of the 18th international conference on World wide web, WWW '09, pp. 761–770. ACM, New York, NY, USA (2009)

[77] Lindstaedt, S., Pammer, V., M¨orzinger, R., Kern, R., M¨ulner, H., Wagner, C.: Recommending tags for pictures based on text, visual content and user context. In: Proceedings of the 2008 Third International Conference on Internet and Web Applications and Services, pp. 506–511. IEEE Computer Society, Washington, DC, USA (2008)

[78] Sigurbjornsson, B., van Zwol, R.: Flickr tag recommendation based on collective knowledge. In: Proceeding of the 17th international conference onWorldWideWeb,WWW'08, pp. 327–336. ACM, New York, NY, USA (2008)