Project no. 34721

TAGora

Semiotic Dynamics in Online Social Communities

http://www.tagora-project.eu

Sixth Framework Programme (FP6)

Future and Emerging Technologies of the Information Society Technologies (IST-FET Priority)

Deliverable D5.5: Impact, Usability and User Communities of web-based experiments

---

Period covered: from 01/06/2008 to 31/08/2009
Start date of project: June 1st, 2006
Due date of deliverable: October 15th, 2009
Distribution: Public

Date of preparation: 15/09/2009
Duration: 39 months
Actual submission date: September 15th, 2009
Status: Final

Project coordinator: Vittorio Loreto
Project coordinator organisation name: PHYS-SAPIENZA
Lead contractor for this deliverable: SONY CSL

Contents

List of Figures

Chapter 1

Introduction

1.1 Purpose of this document

The purpose of this document is to give the reader an overview of the impact, usability and characterisation of the user communities of the web-based experiments performed within the TAGora project. Therefore this document discusses all web-based applications of made by TAGora contractors and points out their respective efforts to improve usability and their respective success in gathering the aimed user basis.

1.2 The importance of impact and usability

Almost all web-based tagging systems, independent of the concept or implementation, are by themselves navigation aiding systems that exploit either the annotating- or the general behaviour of users within the system. For such a system to be useful for its users it is critical to gather a sufficiently large user basis to collect the necessary amount of tagging data to bootstraps the system. To achieve this, the initial system already needs to be useful for single users. Initiating this self amplifying process and tirelessly accelerating it is therefore crucial for any tagging application to be truly successful.

To ensure successful deployment these two aspects of the application or demo need to be addressed with much care. On the one hand it is important to make the application useful for users independent of the availability of tagging data. This means providing a self sufficient benefit which is present even if no tagging data is available at all, to gather an initial user base. On the other hand the system's interface needs to be developed with specific care. It is actually inherently crucial to provide the system with an interface which is simple to learn and intuitive to the user to rise the acceptance of the system to possible users.

The last issue to take into account is to raise the awareness of potential users. This means to inform potential users about the application, through publications or presentations on conferences, blog posts, announcements in universities, etc.

1.3 Document organisation

This document is organised around the individual applications: BibSonomy, Tagster, MyTag, Live Social Semantics, Ikoru and Zexe.net. The NoiseTube application is not covered separately because it is a conceptual extension of Zexe.net and because the system is too recent to evaluate its impact and usability aspects.

Each demo/application is covered in a separate chapter. In general each chapter provides 2 main sections.

The first one points out which impact the application had on users. This includes information about the time and the place at which the application or demo was or still is available to potential users, to which kind of users it was available, how many of them have used or joined the system, about the growth and current size of the user basis and finally about the achievements of user related goals. Furthermore all actions to ensure the dissemination and survival of the application beyond the lifetime of the TAGora project are also covered. Finally this section also provides insight about possible other usages of the application which go beyond the mere functionality of the application or demo itself.

The second section of each chapter is dedicated to the topic of usability. Therein a list of the most important or currently implemented interface features is provided. Each feature is then explained to provide insight into the reasons and purpose of this feature. If done so by the respective partner a short evaluation of the feature is discussed.

Chapter 2

BibSonomy

The BibSonomy system – made by the team at the University of Kassel (UNIK) – is a collaborative resource management system, which supports publications lists as well as bookmark collections. Users who created an account have there the ability to store, organise and share their data with other users. The system makes intensive use of tagging to organize resources and aid navigation.

2.1 Impact

The BibSonomy platform was released even before the start of the TAGora Project and has remained online since then. Over this time the user basis grew constantly, and the collection of valuable data was used to analyse the relations among tags within collaborative tagging systems. Currently the number of registered users run in the thousands, hundreds of which are active on a regular basis. Furthermore, 358 people have subscribed to the mailing list about research issues related to BibSonomy. Access to BibSonomy dataset dumps is also handled via this mailing list. Besides that team also has a blog (http://bibsonomy.blogspot.com) where news about new features and the platform in general is announced.

2.1.1 Publications and press coverage

BibSonomy was covered mostly in publications as a source of collected data used to analyse the semantic structure of tags. BibSonomy was covered in several press articles targeting a broader audience, see http://www.bibsonomy.org/help/contact/press for details.

2.1.2 Integration with other tools and services

BibSonomy has been integrated with a wide range of digital libraries and 3rd party tools and services.

Digital Libraries
Foremost it was integrated into the interfaces several digital libraries as the digital Libraries of the University of Cologne, the institute of Information Sciences at the Saarland University, the KUG library, the Library of the University of Heidelberg and the University of Kassel. Through this integration library users can add publications to their BibSonomy account with one click.

3rd Party Applications
In additional to these direct integrations, the BibSonomy functionality is also available as an optional component for third party software. Through these interfaces BibSonomy data can

easily be integrated into Zope, Moodle or even Wiki- and Webblog Software like XWiki-Page or WordPressBlog.

CiteSmart
Software vendor MireSoft has a product called CiteSmart, a citation management program. It integrates with Microsoft Word and connects with web-based tools like BibSonomy or Connotea. It facilitates importing citation data from these databases and it is able to produce references in various formats for scientific articles written in Word. The BibSonomy teams supports this partnership both to broaden its user base and to make BibSonomy more valuable for existing users.

Zope
Zope[1] is an open source application server for building content management systems, intranets, portals, and custom applications. Publication lists, bookmark lists, and tag clouds, imported from services like BibSonomy, can be dynamically integrated into Zope web pages, using the *KebasData* product.

Typo3 Extension
Typo3 is a popular open-source content management system, used by a large number of private and corporate websites. It offers a generic extension architecture, which enables developers to add custom functionality to Typo3-based websites. For many websites in academic contexts (e.g. personal homepages of researchers, universities, research projects, ...), an important building block is an up-to-date publication list. Maintaining these lists manually is a tedious task. The core concept of the BibSonomy extension for Typo3 is to keep all references cleanly stored inside BibSonomy (leveraging all useful BibSonomy features like import from different formats, scraping services, ...) and to generate publication lists from this data automatically.

Wiki -and Webblog Software
BibSonomy also supports integrations with Wiki- and Webblog Software.

XWiki
To integrate BibSonomy data into an XWiki-page, one has to install the XWiki RSS Feed Plugin XWi (2004). Then, the data can be imported as RSS Feed.

WordPress
Bloggers who are using WordPress can integrate data or components from BibSonomy into their posts – for instance their BibSonomy tag cloud or last three added publications (or all of them). Conversely, blog posts can also be automatically published as bookmarks BibSonomy. To import BibSonomy data to a WordPress blog, the WordPressBlog plug-in [2] has been implemented. A more general way of including BibSonomy content into your system is BibSonomy's JSON feed. JSON (JavaScript Object Notation) is a lightweight data-interchange format, which is now available for all BibSonomy pages.

Moodle
Moodle is a popular e-learning platform for students and lectures Dougiamas (1999). BibSonomy can be integrated to enhance course descriptions and e-learning projects by providing the corresponding literature also via RSS-feeds.[3]

GoogleSonomy
GoogleSonomy is a new extension for Mozilla Firefox which integrates BibSonomy search results directly into Google search results. The extension is customizable so

[1] http://www.zope.org/
[2] http://www.christianschenk.org/projects/wordpress-bibsonomy-plugin
[3] http://educampus.uni-kassel.de

that you can decide whether to search in your personal publications and/or bookmarks, or to search over all BibSonomy posts.

Zotero
BibSonomy now also allows to export citation information to Zotero. Zotero is another Firefox extension, which helps you to collect, manage and cite publications locally in your browser. The other way around (importing data from Zotero into BibSonomy) is not fully automated yet. However, there is a copy and paste workaround.

Scrapers
The collection of scrape tools to automatically extract references from digital libraries or publishers' websites for importing into BibSonomy has been extended, allowing you to store publication metadata automatically from over 60 sites.

Application Programming Interface
BibSonomy now also has an application programming interface (API) which allows external applications to interact with BibSonomy. It is restful [4] API which provides a simple access to all data of BibSonomy. The API documentation is available on http://www.bibsonomy.org/help/doc/api.html. For the Java language, BibSonomy is also offering a client library (available on http://www.bibsonomy.org/help/doc/javaclient.html) One example application which uses the API to access BibSonomy data is JabRef[5], an open source desktop program to manage bibliographic metadata stored in BibTeX databases. The BibSonomy team has extended this tool by using the client part of the API implementation to connect to BibSonomy.

2.1.3 User Basis

During the period of the project the user base has been constantly growing as illustrated in Figure 2.1. This development is a clear indicator that the newly implemented functionalities fulfil their purpose to attract more users to the service.

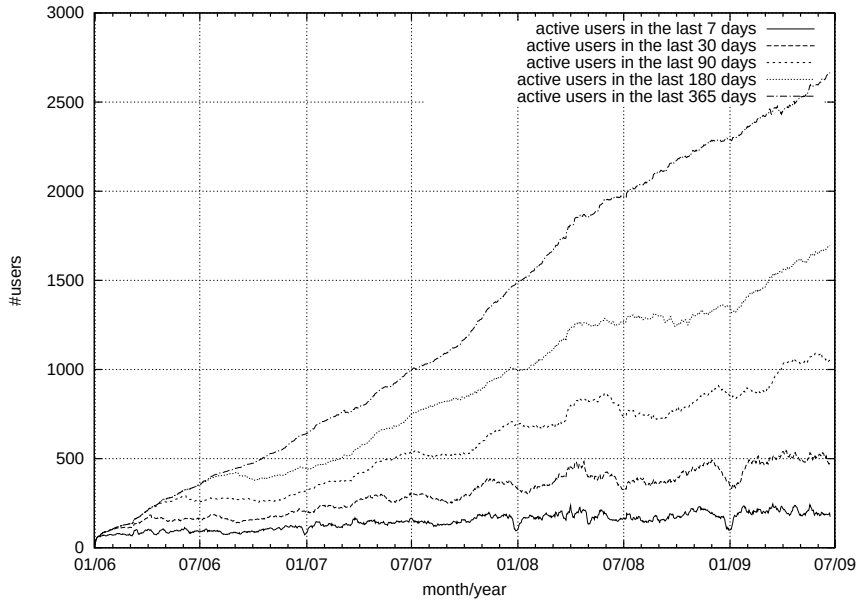

Figure 2.1: Growth of BibSonomy

[4] Using REST (http://en.wikipedia.org/wiki/Representational_State_Transfer) a HTTP-based technique to interact with remote web services.
[5] http://jabref.sourceforge.net

The aim of BibSonomy is to create a huge user community around the web-service. For this purpose users now have the option to create specific groups to organise work groups or collages. Currently the system contains 61 different explicitly defined groups from which most research groups or participants of European projects are.

2.2 Usability

During the TAGora project the development of BibSonomy continued and several new features have been added to the web-application. All these features had the aim to increase the number of subscribed users. These features are described in more detail in Deliverable 2.5.

Chapter 3

Tagster

Tagster – created by the team at the University of Koblenz-Landau (UNI KO-LD) – is a peer-to-peer tagging application. Much like Flickr, Delicious, etc. it allows to tag and share personal data and resources. In contrast to these centralised services where content is uploaded to an Internet service, Tagster follows are decentralised approach in which all content is organised and stored locally on the computers of its users. Tagster is based on a modular architecture, formerly known as the Semantic Exchange architecture (SEA) Franz et al. (2006), which provides the basic functionality for organising and sharing annotated information resources in a decentralised scenario. Additionally, a mechanism for managing distributed tagging statistics is integrated and the application provides different data views for easily navigating the annotated multimedia data.

We have experienced a strong interest in Tagster. Many people liked the idea of sharing their resources in a peer-to-peer fashion, especially with colleagues. However, the attraction of a sufficiently large user group with a long-term interest to use the software is crucial. The client part of the platform should be running for several hours each day. Otherwise, the interest will quickly diminish because nobody can find other people's resources.

3.1 Impact

Different technical problems which were difficult to solve with the available amount of resources, made the attraction of a critical mass of users by the end of the TAgora project questionable. Consequently, the development of Tagster was stopped to divert the effort to other tasks. Still the application was published as open source on Launchpad (https://launchpad.net/tagster) to give other interested people access to it.

3.2 Usability

The aim in the design of the user interface was twofold. On the one hand the same functionality had to be provided as any other centralised folksonomy system has. On the other hand an intuitive way of navigation had to be maintained.

- **Tag-cloud**
  The adoption of typical navigation elements like a tag cloud from centralised systems is motivated by the fact that users of such systems are already accustomed to that type of navigation features. Therefore, the integration of the same or a similar data representations ensures that users get familiar with Tagster quickly.

- **Resource browser**
  The resource browser was implemented in two different ways. Users can either browse

the resources by their associated MIME type. Thus, it is possible to filter a search results such that only images or documents are displayed. The other visualisation type orders all local files in their actual folder hierarchy. Resources from the network are displayed without hierarchical structure since that information is not shared for privacy reasons.

- **Download protocol**
  Since browsing of resources in the network is not enough, Tagster also implements a download protocol for direct retrieval of files from other users. To download a resource through the network the user just has to click on the download button that is displayed next to the tagging data describing the resource. Then the file will be retrieved from the owner and saved in the local download folder. All tags already assigned to the resource will be automatically applied to the downloaded file, too.

Chapter 4

MyTag

MyTag – developed by the team at the University of Koblenz-Landau (UNI KO-LD) with contributions from the team at the University of Southampton (UNI-SOTON) – is a tag-based search engine which allows its users to search different content types like photos, videos and social bookmarks from different sources in parallel. The sources used are popular online resource management and publishing services which rely on tagging, such as Flickr, YouTube and Delicious. MyTag executes search queries transparently and in parallel using the public APIs of these services and presents the retrieved results in separate columns according to content type (images, videos, bookmarks, etc.) as shown by Fig. 4.1.

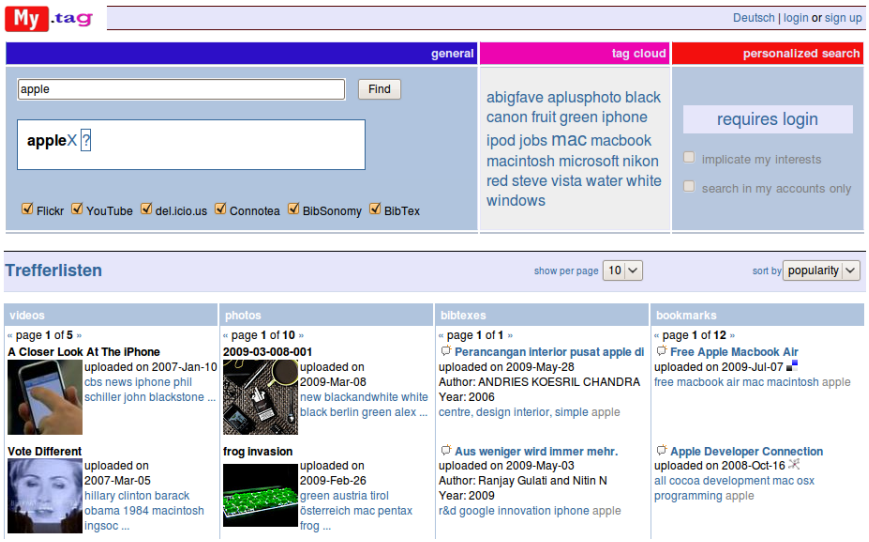

Figure 4.1: Search results as they are presented to MyTag users

While providing search results for the user, MyTag also collects the search interests of registered users and offers them a personalised ranking of results. Additionally, an intelligent search assistant is provided which helps in disambiguating the current search terms by grounding them to possibly relevant articles found in DBPedia [1].

4.1 Impact

MyTag was first released in November 2007. There are several objectives for developing MyTag. First, it should allow its users to perform a cross-media search across different popular tagging

[1] http://dbpedia.org

systems. Second, it should serve as a demonstrator of TAGora technologies.

With regard to the first objective, MyTag was very successful. In its current version, MyTag allows to search in five different tag-based services (BibSonomy, Flickr, YouTube, Delicious and Connotea) and to do a personalized search in these systems, either by restricting the search to ones own resources or by ranking results according to one's own interests.

With regard to the second objective, MyTag was also successful. Shortly after the release of the public BibSonomy API, it integrated BibSonomy into the list of tagging platforms which can be searched with MyTag. Furthermore, it offers an intelligent search assistant which uses the disambiguation web service of Southampton University.

MyTag was also disseminated at several internam occasions to university students and visiting school classes as a demonstrator of TAGora technologies. This led to two bachelor theses which implemented extensions of the original MyTag platform Grabs (2009); Scharek (2009).

4.1.1 Publications and press coverage

MyTag was disseminated at several opportunities. For example, an outcome of the first project seminar was a publication together with the participating students at the WWW 2008 demo and poster session Braun et al. (2008). Furthermore, MyTag was presented during the first Future Internet Symposium Franz et al. (2008) and it was mentioned in the German *Entwickler Magazin* Winkler (2009).

4.1.2 User Basis

In total, 98 user accounts are registered on the MyTag website. Although MyTag offers personalization features to registered users, most of the times MyTag is used by non-registered users. By mid-July 2009, a total of 5200 search queries had been submitted. After dissemination activities like the presentation at the WWW demo and poster session or the presentation at the Future Internet Symposium, the team could observe an increased user activity on MyTag for several days.

4.2 Usability

During the project period the following features have been implemented for MyTag:

- **Integration of Media Sources**
  During the running time of the project five different sources of multimedia content were integrated into MyTag – Flickr, YouTube, Delicious, Connotea and BibSonomy – allowing users to search though all of them at once.
- **Result List Merging**
  By incorporating Bibsonomy and Connotea, search results for bookmarks are the retrieved from three different platforms. This required to introduce an algorithm into Mytag that merges the bookmarks coming from Delicious, Connotea and Bibsonomy into a single result list, to prevent confusing the users. The technical details about the merging algorithm are available in Grabs (2009).
- **Search term disambiguation**
  In a collaborative effort with the Southampton team, an assistant to disambiguate search terms was introduced into MyTag. The current search terms of the user are automatically analysed and sends it to a disambiguation web service offered by the Southampton University. This web service grounds the search tags in articles from DBPedia and returns possible related terms.

MyTag then analyses the returned list of possible meanings of the search terms and filters those meanings which are not represented in the search results retrieved from the tagging platforms. The remaining grounded terms are then presented to the user. The user can then select the intended meaning of the search term and re-rank the current list of results so that resources corresponding to the intended meaning are ranked higher.

The disambiguation service was only introduced in June 2009. Therefore, there currently is not enough data available in order to measure the acceptance of this service by the users of MyTag. However, a user experiment, to analyse in how far this service helps users while search for specific information or resources, is planned for the near future.

- **Searching only through own uploaded contents**
  If provided with the needed account information for all media sources, such as Youtube and Flickr, MyTags offers the ability for the user to search only among the resources he or she uploaded on these platforms personally.

- **Ranking based on "personomy"**
  Based on the "personomy" of a registered user MyTag can adjust the ranking of the search results to match the general personal interest of the users. The personomy is automatically built based on the resources the user picks from the result set. As it is based on the implicit feedback given by selecting from the search results, no additional user effort is required to gain personalization.

- **Multilingual user interface**
  Currently the MyTag user interface is available in English and German language. Further languages might be easily integrated in the future by translating a file with the corresponding text resources.

Chapter 5

TAGnet

People are not fully aware of the metadata they use to annotate resources. Tagging requires little effort and implies no strong committment to consistency: this fosters the externalization of large bodies of metadata, and at the same time makes the structure of the metadata rather unpredictable, even for the user performing the annotation. In the context of a single user, the tag co-occurrence network exposes many of the semantic relations among tags, and between tags and the broader context defined by user's interests and experiences. Visualizing the tag co-occurrence network and allowing the user to manipulate it provides her with a sort of "semantic mirror" that can be used for awareness, for navigation, and for re-organization of metadata. To this end, an application initially not envisioned in the work-plan, *TAGnet*, was designed to exploit the results of WP3 and WP4. The concept was developed by the PHYS-SAPIENZA team in collaboration with the ISI Foundation in Turin. We were able to leverage resources that allowed us to develop this application at no cost for the Consortium. On achieving maturity the application was renamed from *TAGnet* to Netr, because the domain tagnet.org was unavailable for registration. The application was publicly deployed as *Netr* at the URL http://www.netr.it. In the following we will interchangeably refer to this application using any of the two names.

5.1 Impact

TAGnet is a prototype web-based application (http://www.netr.it) designed to provide users with a reflexive tool to expose regularities and patterns in their own tag-based annotations. Tagging patterns can reveal a lot about a user's experience, her interests and her emergent conceptualizations, but users are not aware of these patterns until these regularities are made explicit by means of data analysis and state-of-the-art visualization. TAGnet currently targets Flickr users and provides them with a view on their annotations (tags) that exposes actionable information.

5.2 Perspectives and road-map

In perspective *TAGnet* will be also used as a tool to explore emergent conceptualizations and tag ranking strategies for social annotations. To this end, extensive logging of interface events has been foreseen so that one can compare the measures of node importance and link strength computed by our system with the same notions as explicitly expressed by the user by means of the interface controls. This will yield insights into node ranking and similarity Cattuto et al. (2008a,b); Markines et al. (2009) in folksonomies, as well as a better understanding of what constitutes a better graphical layout (from the perspective of the end-user) in visualizing tag metadata.

In the long term, the user interface of the system will be cloned and customized to set up user studies targeting specific questions on user behavior, emergent categorization and conceptual

structures, as exposed by the annotations of a given user. These experiments will be kept separated from the main system not to impair the applicative goal of *TAGnet*, which will be improved and kept focused as a tool for reflexive exploration of tagging patterns in the context of a single user.

Chapter 6

Live Social Semantics

To disseminate various TAGora technologies at large scale, an experimental application, called Live Social Semantics was designed and developed by the team at the University of Southampton (UNI-SOTON) in collaboration with the SocioPatternsòrg RFID project. The experiment was first deployed at the 2009 European Semantic Web Conference (ESWC09), and subsequently at the ACM Hypertext 2009 (HT2009).

6.1 Approach and goals

Semantic web conferences are a common playground for launching tools and services to enhance the overall social and communicational experiences of attendees. For example, ISWC 2008 used a social networking service to encourage attendees to contact and communicate. ESWC 2006 launched a suite of tools for online presence and chatting. However, such tools are usually outplayed by other more popular systems that provide similar services (e.g. LinkedIn, Facebook, MySpace, Windows Live and Yahoo! messengers). Therefore the experiment launched at ESWC09 was meant to use: (a) the available wealth of linked semantic data, (b) the rich social data from *existing* major Social Networking Systems (SNS), and (c) a physical-presence awareness infrastructure based on active radio-frequency identification (RFID) tags, developed and integrated by the SocioPatternsòrg project of the Institute for Scientific Interchange (ISI) Foundation in Turin, Italy.

Acquiring and integrating these heterogeneous, but overlapping, data sources enabled the team to provide a new experience and services to conference attendees. The main goal was to encourage conference participants to network, to find people with similar interests, to locate their current friends, and to make new ones, using various techniques developed within the TAGora project (e.g. cross-folksonomy analysis, tag filtering and disambiguation, social semantic modelling, tags-to-interest analysis).

6.2 Impact

The Live Social Semantics application was deployed for 4 days (1-4 June 2009) during the European Semantic Web Conference (ESWC'09), which was held in Crete, Greece. More than 300 people attended the conference, out of which 187 accepted to participate in the experiment. Each participant was issued with a uniquely numbered RFID badge. Users were asked to enter their RFID ID number on a website dedicated to this social application. On this website, users were also able to provide their Delicious, Flickr, and Last.FM account information, as well as activating a Facebook application that collected their social contacts.

Out of the 187 people who wore RFID badges, 139 also created accounts in the application web-

Account	Quantity
Facebook	78
Delicious	59
lastFM	57
Flickr	52
Total	246

Table 6.1: Number of social networking accounts entered by Live Social Semantics participants

	Global	Delicious	Flickr
Concepts Generated	1210	922	288
Concepts Removed	247	156	91
Concepts Added	19		
Concepts Saved	982	766	197

Table 6.2: Statistics of the profile generation, editing, and saving.

site. Table 6.1 shows how many social networking accounts were entered into the system by the 139 registered participants.

After the conference, the 49 users who registered on the site but did not enter any social networking accounts were contacted by e-mail. The aim was to understand the reasons behind that. So far 22 responses have been received. Out of those 22 participants, 9 (41%) of them simply did not have *any* social networking account, and only 1 of these 9 indicated that she/he has a Facebook account but has almost no contacts on it. Four participants (18%) indicated that they use *other* networking accounts, (LinkedIn was named twice). Only 2 (9%) of the 22 received replies cited privacy reasons for not sharing their social networking accounts. Six replies (27%) blamed time constraints or the slow Internet connection at the conference venue. One participant (5%) was apparently was too busy.

Out of the 90 people who entered at least one social networking account, 59 of them entered at least one Delicious, Flickr, or Last.FM account (the remaining 31 only entered Facebook accounts, which were not used for generating Profiles Of Interests (POI)). 41 individuals viewed and saved their POI, of which 31 had a non-empty profile generated. Empty profiles were generated for a number of users who registered that had a very small or empty tag-cloud. Table 6.2 summarises the results in terms of the number of concepts automatically generated, the number that were removed manually by users, the number that were added manually, and the size of the final profile they saved.

A total of 1210 DBPedia concepts were proposed (an average of 39 per person across the 31 non-empty profiles), out of which 247 were deleted. While it would be useful to know exactly why users deleted a concept, whether it be simply inaccurate (i.e. incorrect disambiguation), it did not reflect an actual interest (i.e. a very general concept), or it was something they wished to keep private, it was considered too much of a burden to ask users this question when editing their profiles.

The total number of concepts deleted was 20% of those suggested. Although a facility was included on the website for users to add their own interests, few did - only 19 new concepts were added. When comparing the results from Delicious and Flickr, it turns out that 17% of concepts proposed from Delicious tags were deleted, and 32% of those based on Flickr tags. This suggests that the accuracy of topics harvested from Delicious tags was higher than those from Flickr. Inspection of the concepts removed shows that Flickr was likely to suggest concepts referring to years and names.

6.3 Summary

The Live Social Semantics application was a demonstration of how semantics from several different sources can be harnessed and used to enhance the real-world interactions of people at a social gathering. The experiment provided a first opportunity to expose the semantics of social encounters, and investigate recommendation schemes in bodies of data that mix links from social media with links from real-world encounters.

Overall, this application has great potential for further growth over future deployments at conferences and similar events in a wide variety of domains. The experiment was a great dissemination activity for TAGora, and has generated a great amount of interest from many influential organisations and figures, such as STI, EU Commission, Yahoo!, and DERI. The experiment also provided a lot of data that will help to evaluate TAGora technologies.

Chapter 7

Ikoru

Ikoru, the image and music tag-based navigation application, aims to be be an open platform that can be extended with new analysis and visualisation tools. The system was developed at Sony CSL and is primarily used as a test-bed for experimenting with collaborative tagging and content-based analysis. The platform consists of a server-side component and a web interface, which can be seen at http://www.ikoru.net. The content-based tools evolved from a research project that explored the combination of content-based analysis and collaborative tagging. It was initially developed for images but has been extended to handle music files.

On the technological side, we have designed the Ikoru server as a small and stand-alone Web component that can be easily reused and integrated in third-party projects. Running our own servers allows us to gather detailed user data and explore how the analysis of this data can improve tagging systems.

7.1 Impact

We have made the first version of Ikoru available at the end of the first year of the TAGora project. In the second year we have extended the similarity search to audio.

We have kept the web site up and running throughout the TAGora period but have been focusing more on targeted tagging experiments than on the growth of the website in the second half of the project.

This decision had a concrete reason. Despite the strengths of the system, the reality is that in the last two years many sites have integrated tagging features and that these sites can rely on considerable resources and infrastructure to continuously improve their offerings. Technology transfers within Sony have been in principle possible and Ikoru has been presented to many product divisions within the group. However, the collaborations have been not trivial to set up because of the current tendency of Sony to outsource web services.

As a result Sony CSL does not have a precise planning to promote Ikoru to a large audience. Instead, the current strategy is to continue to increase the usability and reliability of the software through its use in small but concrete projects. These focused projects can be managed much more easily and allow us to concentrate on innovative applications of tagging. In the future, we see Ikoru evolve as a generic back-end to store the information about resources, people, and tag assignments. We also see the focus of the tagging applications move away from purely Web-based applications towards real-world applications.

One such project is the artistic installation "Phenotypes/Limited Forms" that was exhibited at the Zentrum für Kunst und Medien (ZKM) in Karlsruhe, Germany, the Bienal de Sao Paulo in São Paulo, Brazil and the "Selective Knowledge" exhibition in Athens, Greece and still is on display in the Museum of Contemporary Art in Siegen, Germany. Although this installation – a joint project

with photographer Armin Linke – is a very particular use of Ikoru, it has allowed us to gather a fair amount of data. More than 8000 visitors picked a selection of eight photos and tagged it using a special "editing table" designed for this purpose. The photos, printed in high-quality on solid boards, are taken from an archive of one thousand photos that are displayed on shelves in the exhibition space. Once the visitors tagged their selection, the editing table prints out a small booklet that they can take home.

To facilitate such small tagging projects by other researchers and developers, and to let Ikoru evolve accordingly, we made the source code available under the GNU Library General Public License (LGPL). It can be found at http://sourceforge.net/projects/ikoru.

7.2 Usability

To allow users to navigation though their image and music files the following features were implemented in the website:

- We showed last year that the image similarity search was a powerful tool, in particular when it is used with tags. Using tags, a visitor can narrow down the number of images that are displayed. For navigation in the content archive, the simple content-based search became a useful tool to select a subset of the images. It can disambiguate, for example, between images of *apple fruit* and *Apple computers*, both tagged with *apple*.

- Recently we enabled Ikoru to store musical melodies Pachet (2008). Compared to photos or audio files, melodies can be analysed and generated at a semantically higher level. It has also the potential to reach a small but passionate community.

- In deliverable D3.3 (Section 3.3) we gave a very brief overview of a thorough study on automatic tag suggestion for music, based on the analysis of the audio signal. In that work, state-of-the-art classification techniques, developed at Sony CSL, were evaluated using a large music database that had been tagged consistently by a group of professionals using a well-defined taxonomy. Although the proposed classification technique performs better than existing techniques, the study also shows that the semantic inference remains extremely difficult. A direct translation of the used method from a clean-room database to an online tagging website is likely to yield unsatisfactory results.

- The implementation of the contextual similarity search for music in Ikoru follows the simple approach that had been successful for images. The features we used are mostly those that are defined in the MPEG-7 audio standard [1]. To test it, we ran the feature extraction algorithms on the Last.fm data set, consisting of more than 18000 tagged music snippets of 26 seconds each. The similarity search can be tested on the Ikoru demo site (http://demo.ikoru.net).

[1] The feature extraction works as follows: the signal is split in overlapping chunks of 2048 samples (approximately. 46 msec long and 23 msec overlap); each chunk is weighted by a Hanning window; for each chunk we apply a DSP operator such as, for example, the root mean square (RMS, related to energy level) or the zero crossing rate (ZCR); we then calculate the first two statistical moments (mean and variance) to aggregate the values of each chunk into a single global value. Most operators return scalar values (RMS, ZCR, ...) except for Mel-frequency cepstrum (MFCC, 20-dimensional) and the Chroma analysis (measures the 12-tone distribution, 12-dimensional). The complete list of operators include: harmonic spectral centroid, harmonic spectral deviation, harmonic spectral spread, pitch, spectral centroid, spectral flatness, spectral spread, spectral kurtosis, spectral skewness, spectral roll-off, ZCR, RMS, RHF, HFC, IQR, centroid, harmonic spectral variation, MFCC, Chroma.

Chapter 8

Zexe.net

The Zexe.net system – developed at the Sony Computer Science Laboratory in Paris (SONY-CSL) – consists of a set of online applications and tools that allow small-scale communities to represent and communicate their views and daily lives on the web. Through the use of smart phones, communities in different cities around the world have published images, videos and sound recordings in Zexe.net for the last five years. Participants not only publish their daily experiences in the form of multimedia files, but they also tag them. Thus, Zexe.net proposes a novel usage for tags by letting users assign them to what we could call "slices of life". We call these web-based tools Community Memories, as they help communities represent and raise awareness about a commons Steels and Tisselli (2008).

8.1 Impact

The Zexe.net platform has been deployed for a range of different communities such as taxi drivers in Mexico City, gypsies in Lleida and León (Spain), prostitutes in Madrid. Two deployments of the Zexe.net system are of particular interest for TAGora: the *canal*MOTOBOY project* and *GENEVE*accessible*. In *canal*MOTOBOY*, motorcycle messengers (called motoboys) in São Paulo, Brazil, report about their journeys. In *GENEVE*accessible* handicapped people in Geneva report about the accessibility of different locations throughout the city.

Although the Zexe.net project already existed in some form since 2003, the applications for the two deployments mentioned here were totally re-written in order to support folksonomies. We found that the concept and the mechanics of tagging were understood very quickly by the participants of these projects, even by those who were not technically literate. The inclusion of folksonomies in these projects greatly improved the way in which the participants dealt with emerging topics and provided a bottom-up way for representing the issues and views of the involved groups in a much more accurate and fine-grained way. By analysing and comparing the tagging activity in *canal*MOTOBOY* and *GENEVE*accessible*, we also show how the scope of a project's focus and tag suggestion can influence the growth and diversity of folksonomies.

Figure 8.1 shows the Tag-Participant network, a tool that shows the participant's position in relation to tags on a 2D plane. Tags are attractors, which means that the closer a participant is to a tag, the more he or she has used it. The objective of this tool is to reflect and compare the participants tagging activities visually.

The following table represents the project's tagging activity figures, as measured on 20/05/2008:

Project	Duration (months)	Users	Tags	Messages	Tag Assignments
*canal*MOTOBOY*	13	15	712	7.975	8.079
*GENEVE*accessible*	3	16	107	2.039	3.188

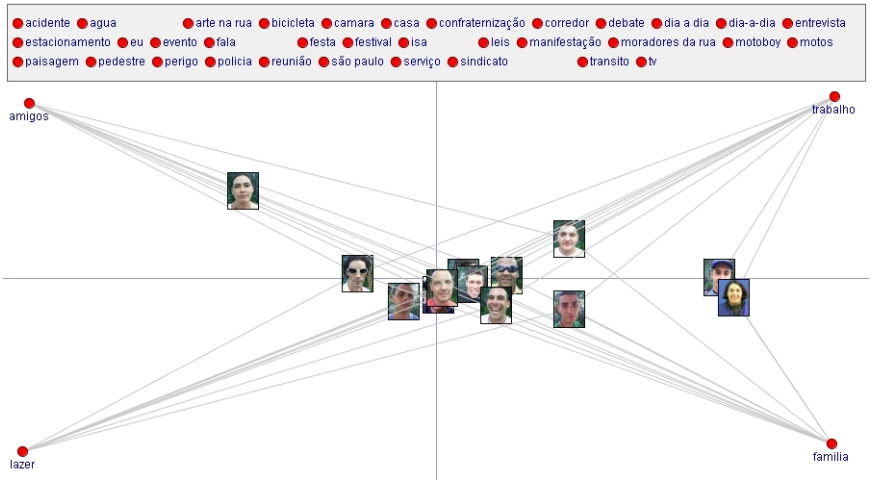

Figure 8.1: Tag-Participant Network in *canal*MOTOBOY*

8.2 Usability

The Zexe.net platform has been deployed for a range of different small scale communities. For each deployment participants were given smart phones and were coached to learn how to use them to report about issues they run into in their daily lives. These reports are aggregated on a web portal where they can be explored by other participants. For navigating these reports on the web portal the following features were implemented:

- Every participant had a GPS-enabled mobile phone. Whenever the GPS is active, geographical coordinates (latitude and longitude) were embedded in every photograph taken (geotagging). Thus, this information is associated to individual multimedia elements. Google Maps is then used to visualise them on a map.
- Multimedia elements, called attachments, are bundled together in *messages*. Messages are sent to the server through the MMS (Multimedia Messaging System) service. Tag assignment occurs on this level, meaning that all the multimedia files included in one message will share the same tag assignment.
- *Messages* are bundled together in larger containers, called *channels*. A *Channel* can belong to a single participant or to multiple ones. Collective channels are created in order to aggregate specific shared topics. Each channel features its own tag cloud.
- Tag clouds can be customised so that they emphasise either the frequency or the popularity of tags. Frequency refers to the number of times a tag has been used, either by a single user (in the case of individual channels) or by the whole group (in collective channels). Popular tags can only be viewed in the tag cloud of the collective channel; their size is proportional to the number of participants who have used them. If a user selects a desired tag *t1* from the tag cloud, the tag is highlighted, and the results are presented. Only co-occurring tags are enabled in the tag cloud for further selection. If the user selects one of these tags, the search is refined by executing a database query that includes all selected tags: *t1* AND *t2* AND *t3* AND ... AND *tn*. Users can deselect any of the previously selected tags at any time.
- Tag clustering was improved by introducing the possibility of manually creating groups of synonyms, which can include any number of tags. For example, the group "marches" can contain its singular and plural forms, "marche" and "marches", and also typographic errors, such as "marhe". A second level of clustering is allowed by the possibility of creating

Figure 8.2: Tag cloud showing highlighted (amigos, familia) and co-occurring (casa, ccsp) tags in *canal*MOTOBOY*

channels from tags. Thus, for example, the channel "obstacles" includes the tag groups "dangers", "déviations", "entrées", "escaliers", etc.

- Keyword searches for tags have been implemented.

- The Zexe.net system also offers an online editing tool which participants use to manage their own channels and perform tasks on individual multimedia elements, such as editing their associated tags or their geographical location. The system also provides a control panel for the system administrator, in which different housekeeping tasks can be performed. Among these tasks are the creation of tag groups, the creation of tag-based channels and the choice of the specific tags to be highlighted on the map.

- For the deployment in **canal*MOTOBOY** we have implemented an automatic detection of the singular and plural forms of the tags. These rules were written specifically for the Portuguese language. Thus, tags which exist in both of these forms are bundled together, publicly displaying only the most popular form.

- In the deployment of **GENEVE*accessible**, the tag search results not only returned a set of photographs, but are showed the corresponding map with the associated locations of those photographs. The map itself also became an interface for navigation through the images, and can be used interactively, together with the tag cloud. In fact, selected tags or tag groups can be associated to markers of a specific color on the map.

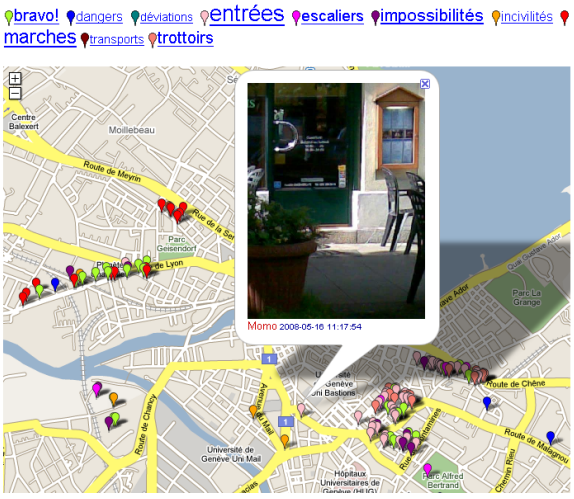

Figure 8.3: The “obstacles” channel in *GENEVE*accessible*, with the tag groups representing the different types of obstacles in the tag cloud. Each tag group has its corresponding marker color on the map.

- A special application installed on the mobile phones provides the users of **GENEVE*accessible** with a list of the 10 most popular tags to choose from when tagging a photograph. This application is constantly updated with data coming from the Zexe.net database. The application also allows users to add a tag which is not found on the list. Because of the specific goals of **GENEVE*accessible**, an initial list of tags to be used for urban obstacles was convened with the participants at the start of the project. Needless to say, these tags quickly became the most popular ones.

Chapter 9

Conclusion

In this last section a brief summery is given of the impact and usability of the applications presented within this document.

BibSonomy
The BibSonomy application continues to extend its user basis steadily while at the same time adding new usability improving features, APIs and integrations with third party library services and tools.

Tagster
Due to technical problems it was difficult the reach a critical user basis for this application. Therefore the further development of the application was stopped. However, the last version of the application is fully functional and is still available for potential users. Furthermore, the source code was published under an open source license.

MyTag
The development of the MyTag application as a demonstration was completed and allowed a range of students to work with TAGora related topics. Additionally the application provided topics for two bachelor thesis.

TAGnet
TAGnet is a prototype web-bases application designed to provide users with a reflexive tool to expose regularities and patterns in their own tag-based annotations. TAGnet currently targets Flickr users and provides them with a view on their annotations (tags) that exposes actionable information.

Live Social Semantics
The application was disseminated on two conferences as a demonstration and gathered a considerable user basis. This allowed a deep insight into the general and connection of users to their various social networks.

Ikoru
Despite expectations the team at SONY-CSL realised that they did not have the resources or the support from within SONY to promote Ikoru to a bigger audience, and therefore decided to focus the system on small scale applications. It was highly successfully deployed as part of the "Phenotype / Limited Forms" installation of artist Armin Linke. Furthermore the development of new usability features has not stopped. To disseminate the application further the source code was released under an open source license.

Zexe.net
The Zexe.net platform wars successfully deployed for several communities. The inclusion

of folksonomies in the application greatly improved the way in which the participants dealt with representing the issues and views of the involved groups. By analysing and comparing the tagging activity in *canal*MOTOBOY* and *GENEVE*accessible*, the application has also shown how the scope of a project's focus and tag suggestion can influence the growth and diversity of folksonomies. Within the scope of TAGora, the ideas behind the Zexe.net platform have lived on in its successor NoiseTube which extends its concepts to create a collaborative community platform to sense and tag exposure to pollution.

Bibliography

XWiki. 2004. http://www.xwiki.com.

Max Braun, Klaas Dellschaft, Thomas Franz, Dominik Hering, Peter Jungen, Hagen Metzler, Eugen Müller, Alexander Rostilov, and Carsten Saathoff. Personalized Search and Exploration with MyTag. In *Proceedings of the WWW 2008 Poster Session*, 2008.

Ciro Cattuto, Dominik Benz, Andreas Hotho, and Gerd Stumme. Semantic Analysis of Tag Similarity Measures in Collaborative Tagging Systems. In *Proceedings of the 3rd Workshop on Ontology Learning and Population (OLP3)*, pages 39–43, Patras, Greece, July 2008a. ISBN 978-960-89282-6-8. URL http://olp.dfki.de/olp3/. ISBN 978-960-89282-6-8.

Ciro Cattuto, Dominik Benz, Andreas Hotho, and Gerd Stumme. Semantic Grounding of Tag Relatedness in Social Bookmarking Systems. In Amit P. Sheth, Steffen Staab, Mike Dean, Massimo Paolucci, Diana Maynard, Timothy W. Finin, and Krishnaprasad Thirunarayan, editors, *The Semantic Web – ISWC 2008, Proc.Intl. Semantic Web Conference 2008*, volume 5318 of *LNCS*, pages 615–631, Heidelberg, 2008b. Springer. URL http://dx.doi.org/10.1007/978-3-540-88564-1_39.

Martin Dougiamas. Moodle. 1999. http://www.moodle.de.

Thomas Franz, Carsten Saathoff, Olaf Görlitz, Christoph Ringelstein, and Steffen Staab. SEA: A Lightweight and Extensible Semantic Exchange Architecture. In *Proceedings of the 2nd Workshop on Innovations in Web Infrastructure. 15th International World Wide Web Conference (Edinburgh, Scotland)*, 2006.

Thomas Franz, Klaas Dellschaft, and Steffen Staab. Unlock your Data: The Case of MyTag. In *Proceedings of FIS-2008 - the Future Internet Symposium*, Vienna, Austria, 2008. Springer.

Daniel Grabs. Beschreibung und Evaluation des MyTag Merge Algorithmus. Master's thesis, Universität Koblenz-Landau, 2009.

Benjamin Markines, Ciro Cattuto, Filippo Menczer, Dominik Benz, Andreas Hotho, and Gerd Stumme. Evaluating Similarity Measures for Emergent Semantics of Social Tagging. In *18th International World Wide Web Conference*, pages 641–641, April 2009. URL http://www2009.eprints.org/65/.

F. Pachet. Description-Based Design of Monophonic melodies, 2008. Submitted to Computer Music Journal.

Matthias Scharek. Optimierung von Suchmaschinen basierend auf dem Suchverhalten von Benutzern im Internet. Master's thesis, Universität Koblenz-Landau, 2009.

L. Steels and E. Tisselli. Social Tagging in Community Memories. In *Proceedings of the 2008 AAAI Spring Symposium*, Social Information Processing, Stanford University, California, USA, 2008.

Reimar Winkler. Ein zweiter Boden für das Internet. *Entwickler Magazin*, pages 104–105, January 2009.