

Semantics, Sensors, and the Social Web: The Live Social Semantics Experiments

Martin Szomszor¹, Ciro Cattuto², Wouter Van den Broeck²,
Alain Barrat^{2,3}, and Harith Alani⁴

¹ City eHealth Research Centre, City University London, UK

² Complex Networks and Systems Group

Institute for Scientific Interchange (ISI) Foundation, Turin, Italy

³ Centre de Physique Théorique (CNRS UMR 6207), Marseille, France

⁴ Knowledge Media Institute, The Open University, UK

Abstract. The Live Social Semantics is an innovative application that encourages and guides social networking between researchers at conferences and similar events. The application integrates data from the Semantic Web, online social networks, and a face-to-face contact sensing platform. It helps researchers to find like-minded and influential researchers, to identify and meet people in their community of practice, and to capture and later retrace their real-world networking activities. The application was successfully deployed at two international conferences, attracting more than 300 users in total. This paper describes the Live Social Semantics application, with a focus on how data from Web 2.0 sources can be used to automatically generate Profiles of Interest. We evaluate and discuss the results of its two deployments, assessing the accuracy of profiles generated, the willingness to link to external social networking sites, and the feedback given through user questionnaires.

1 Introduction

Most conference attendees would agree that networking is a crucial component of their conference activities. It is vital for researchers to network and foster collaboration, with many ties originating in casual meetings and spontaneous conversations. Networking at conferences can be very daunting, especially to junior researchers and those who cross discipline boundaries. Furthermore, researchers often lose track of whom they met and where at such events. Such problems become even more evident in medium to large size conferences, where it is easy for individual researchers to get lost in the crowd. Unfortunately, there are few applications to help researchers to initiate, capture, and preserve their online as well as offline social interactions during conferences.

To this end, we have developed Live Social Semantics (LSS), a novel application that brings together data from the Semantic Web, Web 2.0, and a face-to-face (F2F) contacting sensing platform. LSS encourages and strengthens collaboration and communication between researchers by supporting their social networking activities during conferences, helping them find and locate like-minded individuals, people in their community of practice, and to permanently log, and retrace, their F2F contacts. LSS integrates (a) academic related data from the Semantic Web, (b) the rich social data

from major social networking sites, and (c) a physical-presence awareness infrastructure based on active radio-frequency identification (RFID). This application was deployed at two major international conferences, ESWC 2009 in Crete and HyperText 2009 in Turin, where it was used by 300 researchers with promising results.

The next section describes a variety of related works, followed by a full description of the LSS architecture in Section 3. Section 4 presents our method for building and representing Profiles of Interest. Results and evaluation of LSS deployments are covered in Sections 5 and 6 respectively. Discussion and future work is given in Section 7, followed by conclusions in Section 8.

2 Related Work

The interplay of networking and social contact at a conference gathering was initially investigated in the context of opportunistic networking for mobile devices [12] by using wearable Bluetooth-enabled devices. Subsequent work focused on sensing organisational aspects [6] by using Bluetooth-enabled mobile phones, and on characterising some statistical properties of human mobility and contact [21,16]. These early experiments could not assess face-to-face (F2F) human contact in a large-scale setting since they mostly relied on Bluetooth communication. Wu and colleagues used what they call “sociometric badges” to investigate the impact of F2F interactions on productivity [22]. These badges used radio frequency to detect physical proximity, infrared to detect F2F body alignments, and voice sensors to detect conversations.

RFID is an increasingly popular technology for location tracking. IBM used RFIDs to track attendees of a conference in Las Vegas in 2007. The devices were used to track session and meal attendance [21]. The information they collected was limited to the name, title and institution of attendees. No social or semantic data was collected or used. Fire Eagle¹ by Yahoo! is a service that detects the geographical location of users (e.g. based on WIFI access points), and allows them to share it with their online friends.

Recently, the SocioPatterns project² investigated patterns of human contact at large-scale social gatherings by deploying a distributed RFID platform that is scalable and attains reliable detection of F2F interactions as a proxy of social contact [3]. The LSS application presented here leveraged that platform to mine real-time social contacts.

To the best of our knowledge, our LSS application is the first where real-world F2F contacts are mashed up in real time with semantic data from online tagging systems. The free nature of tagging generates various vocabulary problems: tags can be too personalised; made of compound words; mix plural and singular terms; they can be improper words; they can be synonymous, etc. [14,9,10]. This total lack of control obstructs analysis [13]. In our work, we follow the approach of *cleaning* existing tags using a number of term filtering processes, similar in spirit to those used in [11].

LSS constructs semantic models of interests for individuals by merging and processing their tagging activities from multiple online social networking systems (SNS). This process involves dealing with several problems, such as filtering of tags, disambiguating them, associating tags with semantics, and identifying interests. Tag ambiguity is a well

¹ <http://fireeagle.yahoo.net/>

² <http://www.sociopatterns.org>

recognised problem in folksonomies. Clustering has been investigated as a disambiguation approach, where tags are grouped together based on their co-occurrence, to facilitate distinction between their different meanings [4,23,17]. While such techniques have demonstrated that the underlying folksonomy structure does contain information that can enable automatic disambiguation, they are too computationally expensive and lack any semantic grounding. Angeletou and colleagues [2] used WordNet to identify ambiguous tags, and compared the WordNet senses for the tag to those of the co-occurring tags, to identify the most similar sense. In our approach, we used DBpedia³ for disambiguating tags, automatically associating them with Semantic Web URIs. Some manually-driven approaches have been proposed for assigning URIs to tags (e.g. [17,15]). Similarly to [20], we explore a strategy for the automatic selection of URIs using DBpedia concepts [7].

3 Live Social Semantics Application

3.1 General Architecture

The system architecture of LSS is shown in Figure 1. The diagram is vertically partitioned into two spaces: the online world (i.e. data about individuals held on the web), and the physical space (i.e. RFID-based contact data). Data in the online world is sourced from the following:

- Social networking sites: Tagging and social relation data is collected from Delicious, Flickr, Facebook, and lastFM using the Extractor Daemon. This data is then used to reflect the online contact network of individuals. The tagging data is processed by the Profile Builder (centre, top of diagram) to infer their interests. The Tagora Sense Repository is responsible for associating tags to URIs from DBpedia.
- Semantic Web Linked Data: Information on publications, projects, and the Community of Practice (COP) of researchers is retrieved from RKBExplorer⁴ [8] and semanticweb.org. This data is used to reflect the contact network of individuals based on their paper co-authorships and project co-memberships.

Physical space data is collected from F2F contacts between individuals which are measured using RFID readings (Section 3.2). Such data is fully integrated with the online world data in a triple store (centre right of Figure 1), where all the data is stored. This enriches the visualisation and processing of real-world social contacts with the online social contacts of those individuals. A focused *Contact* ontology⁵ was used to represent real-world social interactions between individuals, recording the total contact time on a daily basis (sections 3.2 and 3.3).

3.2 Real-Time Social Contacts

Real-world interactions of conference attendees are mined using RFID hardware and software infrastructure developed by the SocioPatterns project [3]. Willing participants

³ <http://dbpedia.org/>

⁴ <http://www.rkbexplorer.com>

⁵ <http://tagora.ecs.soton.ac.uk/schemas/LiveSocialSemantics>




Fig. 1. Live Social Semantics Architecture

were issued with RFID badges to monitor their F2F contacts with others. The RFID badges engage in multi-channel bi-directional radio communication, and by exchanging low-power signals which are shielded by the human body, they can reliably assess the continued F2F proximity of two individuals. A F2F contact is recorded if users face each other for around 20 seconds or more, within a distance of around one meter.

We generate a weighted graph to represent the cumulative F2F contacts between the participants. This information is periodically uploaded to the triple store via RDF/HTTP and integrated with the other data layers.

We use the real-world and online social relations to compute simple recommendations. For example, if two attendees are in F2F contact at a given time, the server searches for, and displays, any mutual contacts from the online world data, for example people who are not present at the given time, but are nevertheless connected to the two users in one of the online social networks used by LSS (Section 3.3). Details of using RFID in LSS can be found in [5,1].

3.3 Visualisation

LSS has two visualisations, taken from the SocioPatterns project (detailed in [1]):

- **Spatial View:** This view provides an overview of the real-time contact graph. It represents the location of RFID-badge wearing participants within range of the RFID readers, as well as their ongoing social contacts. Each participant is represented by a labelled yellow disc or, when available, by their Facebook profile picture. The contacts are represented by yellow edges, whose thickness and opacity reflects the weight of the contact. The edges are decorated, where applicable, with small

Facebook, Flickr, Delicious, lastFM or COP icons, marking the occurrence of that relationship in the respective network.

- User-focus view: This view displays the social neighbourhood of a particular user. It shows all participants with whom this user has ongoing contact with (yellow edges for live contacts) or had significant (cumulative) contact with (grey edges for historical contacts). This view also attempts to *close relevant triangles*, by showing mutual contacts as explained earlier.

4 Semantic Profiles of Interest

Tags usually reflect the interests of their authors. Such interests could range from topics, places, events, people, hobbies, etc. We have developed a tool that processes the public tagging activities of users and automatically generates a list of DBpedia URIs to represent the interests of the taggers [18]. To generate Profiles of Interest (POI) from social tagging, we follow these steps (described in the following sections):

1. **Collect tagging information:** A user's complete tagging history is extracted from the target site using public APIs or screen scraping, and converted to RDF. We utilise our Tagging Ontology to represent all tagging events, recording the resource tagged, the tag ordering, and date of annotation.
2. **Associate Tags with Potential Concepts:** Using the TAGora Sense Repository, we associate each tag to a set of potential DBpedia URIs that represent the tag senses.
3. **Perform Tag Disambiguation:** For tags with more than one candidate sense, we perform some basic disambiguation to discover the intended meaning.
4. **Calculate Interest Weights:** For each DBpedia URI identified as a potential interest, we calculate a weight based on tag frequency and a time decay factor.
5. **Create Profile of Interest:** Generate a ranked list of interests.
6. **User Verification:** Allow users to verify and edit their POI as they see fit.

4.1 Collecting Tagging Data

The first step in building a POI is to collect social tagging information from various folksonomies. In previous work we used the Google Social API⁶ to find and correlate several social networking accounts of given users [19]. In LSS however, users must explicitly enter their social networking accounts on the LSS website. Therefore users are given full control in deciding which of their accounts will be used and shared.

The data collection process is responsible for harvesting information from a range of social networking sites. In the case of Flickr, Facebook, and Last.fm, APIs are provided that allow us to download a complete history of user tagging activity. However, for Delicious, the API is limited and so we used custom screen-scraping scripts. All tagging information is stored by LSS in RDF. We use the TAGora tagging ontology⁷ (Figure 2) which is specifically designed to represent tag assignments (posts) and tag use (frequency, time, relation to Global Tag, etc). In future work we plan to merge this ontology

⁶ <http://code.google.com/apis/socialgraph/>

⁷ <http://tagora.ecs.soton.ac.uk/schemas/tagging>

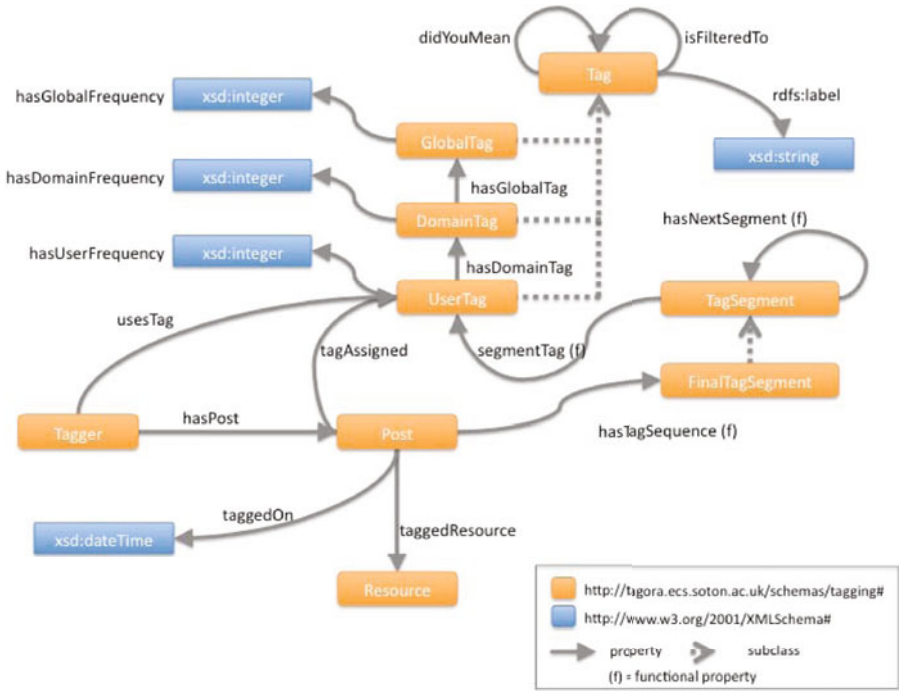


Fig. 2. TAGora Tagging Ontology

with SIOC,⁸ by extending the latter to include our detailed tagging representations that are necessary for tag disambiguation.

4.2 Associate Tags with Potential Concepts

Tags can be misspelled, synonymous and come in a morphologic variety. As a result, important correlations between resources and users are sometimes lost simply because of the syntactic mismatches of the tags they used. To this end, we developed the TAGora Sense Repository⁹ (TSR), a Linked Data enabled service endpoint that provides tag filtering services and extensive metadata about tags and their possible senses. When the TSR is queried with a particular tag string, by formatting a URI that contains the tag (in a REST style, e.g. <http://tagora.ecs.soton.ac.uk/tag/apple/rdf>), the tag is filtered, matched against a set of potential DBpedia and W3C Wordnet URIs that represent possible meanings for the tag. For the purposes of the LSS application, we also provide a SPARQL endpoint. In the following Sections, we briefly describe the functionality of the TSR in terms of the index we built and the search API provided.

Creating The Resource Index. The first stage in building the TSR was to process the XML dump of all Wikipedia pages to index all titles, mine redirection and disambiguation links, and extract term frequencies for each of the pages. For the current version

⁸ <http://sioc-project.org/sioc>

⁹ <http://tagora.ecs.soton.ac.uk/tsr/>

we use a dump available from <http://download.wikimedia.org>, created on 08/10/2008. For each Wikipedia page in the dump, we extract and index the page title, a lower case version of the title, and a concatenated version of the title (i.e. the title `Second.Life` becomes `secondlife`). This multiple title indexing enables us to match more easily tags that are made up of compound terms. We also extract redirection links, disambiguation links, as well as the terms contained in the page and their frequencies. During this indexing process, we also store a list (and count) of all incoming links to each page. Since the dump is large, we only store terms with a frequency greater than the mean frequency of all terms in that page. This data is stored in a triple store using our own extended DBpedia ontology since we are providing more detailed metadata about the entries than DBpedia.org, such as the term frequencies. Each Wikipedia page in the TSR is also linked to DBpedia via the `owl:sameAs` property.

Searching For Senses. When the TSR is queried with a tag, the first step is to find a list of candidate DBpedia resources that represent possible senses of the tag. We begin by normalizing the tag string (i.e. removing non-alphanumeric characters as described in [12]). The triple store is then queried for all entries with the same lowercase title or concatenated title as the tag. During this process, we are likely to encounter redirection links and/or disambiguation links, both of which are followed. When a set of candidate senses has been created, we calculate the total number of incoming links for each resource (including the sum of incoming links for any pages that redirect to it). Finally, a weight is associated with each possible sense, calculated by dividing the number of incoming links associated with that sense by the total number of incoming links for all senses associated with the tag. This basic page rank inspired measure means senses that have very specific meanings receive much lower weights than those associated with general concepts.

For each user tag in LSS, we use a property in the Tagging ontology that links it to the Global Tag in the TSR. Figure 3 shows how a FOAF profile for an LSS user (denoted with the URI `tagora:eswc2009/foaf/4`) is linked to a representation of their Delicious activity (with the URI `tagora:delicious/martinszomszor`). The property `tagging:usesTag` links their FOAF URI to each of their Delicious Tag URIs (e.g. `ontologymapping`) that is, in turn, linked to the Global Tag URI in the TSR (`tagora:tag/ontologymapping` in this case). The TSR provides a link from the Global Tag to the possible DBpedia senses (via the `disam:hasPossibleSense` property). These links can be used to infer that an individual is potentially interested in a particular concept, in this example the tag `ontologymapping` is mapped to the DBpedia entry for `Semantic_Integration`.

4.3 Tag Disambiguation

The disambiguation process aims to analyse the context in which a tag has been used to identify the most likely sense among all possible senses for that tag. Tags are considered ambiguous if they are associated with multiple senses (i.e. more than 1 DBpedia resource). For such a tag, its context is captured and represented as a term vector. By context we mean the other tags that were used to annotate a given resource, hence each use of the tag can have different contexts. We construct another vector from term

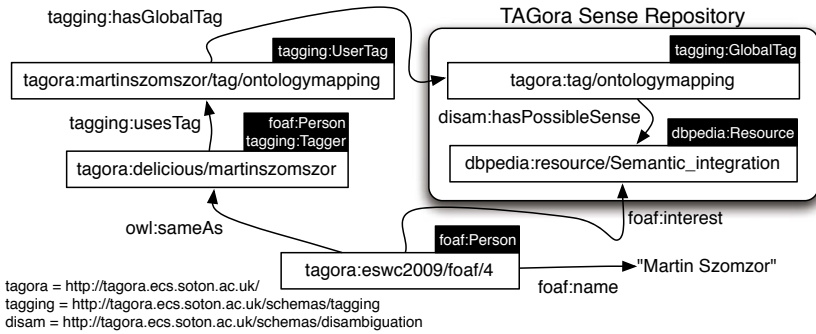


Fig. 3. Linking users to interests inferred from their tagging activities on social networking sites

frequencies associated with the possible DBpedia senses. We then measure cosine similarity between these vectors, and if one of the similarity scores is above a threshold (0.3 in this case), we conclude that this is the correct sense for that tag. If more than one (or zero) senses score above the threshold, we do not associate a meaning to the tag since we cannot reliably choose a correct sense. By iterating through all tags associated with a user (i.e. through Delicious or Flickr), we are able to build a candidate resource list of interests C . Details of our disambiguation algorithm and some initial experiments can be found in [7].

4.4 Calculating Interest Weights

For each interest (i.e. DBpedia resource) $r \in C$, we calculate a weight $w = f_r * u_r$, where f_r is the total frequency of all tags disambiguated to sense r , and u_r is a time decay factor. Therefore, tags that have been used more recently will receive a higher weight than those used earlier in time. If many tags of a given user are associated with the same interest, then the weight for that interest will increase accordingly. The final list of interests contains only those with a weight above the average weight for that user.

4.5 Creating the Profile of Interest

The lists of interests produced by the previous processes are used to generate an RDF Profile of Interest (POI) for each users using the FOAF `interest` property to link the person to the relevant Wikipedia categories. If more than 50 candidate interests have been found, we rank them by weight and suggest the top 50.

4.6 User Verification

Once POIs are generated, the users can browse the list of interests and edit as required, removing or adding new interests as they see fit (Figure 4). Users may wish to remove an interest for various reasons, for example if it was incorrectly identified (e.g wrong disambiguation or filtering), if it is not an interest (or a historical one), or if the user chooses not to share it with the community (private or deemed irrelevant). Users can authorise LSS to use their profiles by clicking on a *save* button.

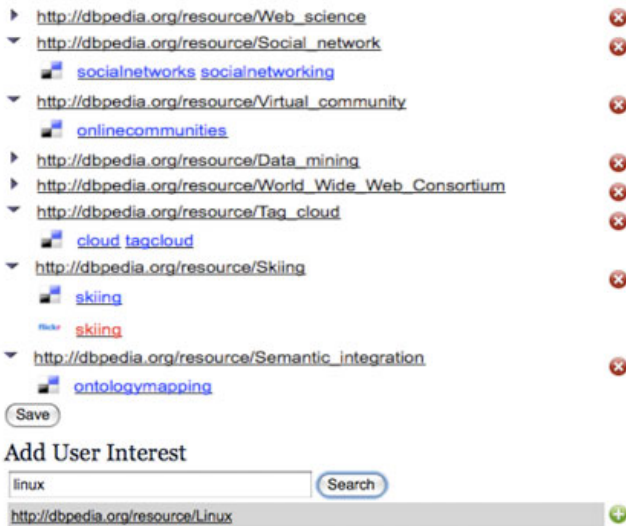


Fig. 4. Users can browse and edit their profiles of interest before authorising their use. The tags and their sources that led to each interest are shown. These profiles are automatically generated from users' public tagging activities.

5 Experiments and Results

The LSS application was deployed for a total of 7 days at the following two events:

- European Semantic Web Conference (ESWC) in Crete, 1-4 June 2009: This conference was attended by 305 people, out of which 187 participated in LSS. Out of the 187 who collected an RFID badge, 139 of them also created accounts on our application site.
- HyperText (HT), Turin, June 29-July 1, 2009: Attended by around 150 people. 113 of them collected an RFID, and 97 registered with LSS.

Each participant was issued a uniquely numbered RFID badge. Users were asked to enter their RFID ID number on the dedicated LSS website. On this website, users were also able to provide their Delicious, Flickr, and lastFM account names, as well as activating a Facebook application that collected their social contacts. The results reported below focus on user participation and SNS account declaration and POI generation. Results and statistics of RFID use can be found in [5].

Participation results. As mentioned above, out of a total of 455 attendees of the ESWC and HT conferences, 300 of them took part in Live Social Semantics (187 at ESWC and 113 at HT). Out of these 300 users, 236 of them created an account on the application site (139 at ESWC and 97 at HT). Hence around 21% of the users who collected an RFID badge did not register to submit any information about themselves (e.g. name, email, social network accounts). F2F contacts of such users were captured, but were not associated with any personal profile.

Table 1. Number of social networking accounts entered into LSS by 236 users during two field experiments

Experiment \ Account	Facebook	Delicious	lastFM	Flickr	Total
	ESWC09	78	59	57	
HT09	48	28	26	23	125
Total	126	87	83	75	371

Table 2. Number of users who entered 0, 1, 2, 3 or 4 social networking accounts into the Live Social Semantics site during experiments at ESWC09 and HT09 conferences

Experiment \ Number of accounts	0	1	2	3	4	Total
	ESWC09 users	49	36	28	13	
HT09 users	35	18	23	8	13	97
Total	84	54	51	21	26	236

Table 3. Table shows the number of interests generated from tags taken from Delicious, Flickr, or lastFM, and how many were removed by users. These statistics are based on 72 POIs verified and saved by their owners.

	Global	Delicious	Flickr	lastFM
Concepts Generated	2114	1615	456	43
Concepts Removed	449 (21%)	307 (19%)	133 (29%)	9 (21%)

Declaration of social networking accounts. Users were able to declare on the LSS site their accounts for Delicious, Flickr, lastFM, and Facebook. The numbers of such accounts that our 236 registered users (i.e. users with LSS accounts) declared on the LSS site are shown in Table 1. The number of social networking accounts declared on LSS site by each individual user varied from 0 (i.e. did not enter any accounts), to 4 (i.e. entered an account for each of Delicious, Flickr, lastFM, and Facebook). Table 2 shows that about 36% of our 236 registered users did not declare any social networking accounts. It also shows that around 58% of our users declared more than one social networking account.

Semantic Profiles-of-Interest results. We analysed 72 POIs that were verified and activated by our users (Section 4.6). Table 3 shows the total number of interests that were automatically generated, and those that were removed manually by users during both field experiments. A total of 2114 DBpedia concepts were proposed, out of which 449 were removed by users (21%). Although a facility was included on the website for users to add new interests, only 19 new concepts were added.

6 Evaluation

Table 3 above showed that 29% of interests suggested from Flickr tags were removed by users, in comparison to 19% and 21% for Delicious and lastFM respectively. This suggests that Delicious and lastFM are perhaps more reliable sources of user interests

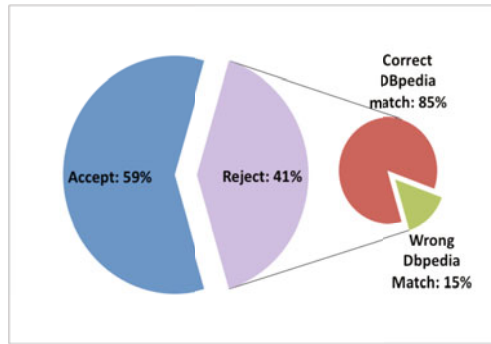


Fig. 5. 11 users edited their POIs in HT09. They accepted 59% of the interests that our system generated, and rejected 41%, out of which 15% were matched by our system to incorrect DBpedia URIs (6% of all suggested interests).

than Flickr. Inspection of the concepts removed shows that Flickr was likely to suggest concepts referring to years, names, or to places that users visited in the past, but are no longer interested in.

To evaluate the accuracy of our interest suggestions, we examined the interests that our users removed from their profiles during the HT09 experiment¹⁰. Users may choose to delete an interest because it is simply inaccurate (i.e. wrong DBpedia match), it does not reflect an actual interest (i.e. is a very general concept), or it is something they prefer to keep private. Users seem to have different perceptions of what an interest is, or which ones are worthy of sharing in this context. Some users were very conservative and only kept a few of the interests that our system generated for them, while others kept almost all their proposed interests. In future LSS implementations we intend to allow users to instantly input their rationale for removing an interest. Understanding these drivers will help us to better design and tune the POI generation process. However, for this evaluation, we will focus on finding out how many of the removed interests were based on tags that our POI process matched to irrelevant DBpedia URIs.


Although 36 of our users at HT09 activated their POIs (by saving them - Section 4.6), only 11 of them removed any interests. The other 25 users might have been totally satisfied with their original POIs, or perhaps they saved their profiles without reviewing them. To be on the safe side, in this evaluation we focus on the POIs that were clearly *scrutinised* and *corrected* by their owners. On average, those 11 HT09 users kept 63%, 57%, and 49% of the interests that the system suggested based on their Delicious, Flickr, and lastFM tagging activities. Several users removed lastFM-based interests, although those interests referred to the music bands that these users listened to the most.

Figure 5 shows the percentages of tags that these 11 users removed from their automatically generated POIs. Although these users removed 41% of their POIs, only 15% (30 tags out of 203 removed ones) of these removed tags were given the wrong

¹⁰ Profiles of interests from ESWC 2009 were later anonymised and hence could not be included in this detailed evaluation since it requires an examination of users' original tags and tagged resources.

Table 4. Reasons why some users didn't enter any social network accounts to our application site

Option	Reason	No. Users	%
a	don't have those accounts (or rarely use them)	16	44%
b	use different networking sites	10	28%
c	don't like to share them	3	8%
d	didn't get a chance to share them (e.g. no computer, slow Internet)	6	17%
e	other	1	3%
Total		36	100%

**Fig. 6.** Comparison of users' answers to our participation survey for the two experiments

semantics (i.e. matched to the wrong DBpedia URIs). For example, for most users the tag “km” was wrongly matched to the concept “Kilometre” in DBpedia, instead of “Knowledge Management”. With a closer look at those 15% of tags, we find that 2% of them originated from Flickr, and the rest came from Delicious. This is hardly surprising since Delicious tags tend to be more diverse than those from Flickr. The majority of Flickr tags referred to known geographical places that have dedicated DBpedia URIs.

In addition to the above, we have also evaluated the shareability of SNS accounts by our users. As mentioned in Section 5, a total of 84 registered users did not enter any social networking accounts on the LSS site. To understand the drivers behind this, we ran a survey where we asked each of these users to pick their main reason out of the 5 options shown in Table 4. We received 36 responses to our survey so far, and the main reasons the users picked are listed in Table 4. It is clear that not having any SNS account is the most common reason for not declaring any. LinkedIn¹¹ and xing¹² were mentioned by several users as alternative SNS accounts, which LSS does not yet support. Although several users mentioned privacy concerns, only 8% of the users selected this as their primary reason.

Figure 6 shows a comparison of the answers to our survey (Table 4) from the ESWC and HT experiments. It is interesting to see that answer *d* was very common for ESWC attendees, who often blamed the unreliable Internet connection at the venue for their inactive participation, whereas this was not an issue for HT.

¹¹ <http://www.linkedin.com>

¹² <http://www.xing.com>

7 Discussion and Future Work

The first phase of LSS development, which led to the prototype tested at ESWC09 and HT09, was focused on architectural design and technology integration to demonstrate a novel proof-of-concept application. The second phase of development will focus on scalability, extendibility, and services.

The LSS application has so far only supported 4 currently popular SNSs: Delicious, Flickr, lastFM, and Facebook. We plan to extend LSS to allow users to submit their FOAF files, and to support other networking sites (e.g. Twitter, LinkedIn, xing). We also plan to develop an open plug-in architecture to allow external parties to develop connection to other networking systems to LSS.

The generation of profiles of interests from social tagging systems produced promising results. These POIs highlighted various general interests of users that usually cannot be inferred from their publications or project descriptions (e.g. “skiing”, “iPhone”, “sewing”, “The Beatles”). However, many users did not take the crucial step of verifying and editing these profiles. This might be due to a misunderstanding of the purpose or value for taking this step. We hope that more users will be encouraged to edit their profiles once we provide additional services that use these profiles, for example, to highlight people with similar interests.

Users who saved their profiles during both experiments removed approximately 21% of the interests the system suggested. This goes up to 41% if we exclude users who saved their profiles without any modifications. Although there could be many reasons for why users choose to remove an interest, our investigation showed that only 15% of the removed interests were totally inaccurate tag-to-concept associations (Section 6). The other 85% were proper associations, but did not necessarily represent an interest. This is a clear indication that we need to develop more sophisticated methods for determining what constitutes an interest and what does not. One promising approach is to tap into our users’ collective intelligence to improve our POI generation process, for example by filtering out the interests that most users tend to reject (“Tutorial”, “API”) or those that are too common or too general (e.g. “web 2.0”, “Semantic Web”).

Next step for interest identification will be to model users’ interests in semantic hierarchies, enabling us to represent interests at different levels of granularity. For example, if someone is interested in “Visualbasic”, “Perl”, and “C++”, then one can infer that this person is interested in “Programming languages”. The hierarchy can show how general the user interest is, so one user may use the tag “music” very often, while another might tag with “jazz” or “Hip hop”, which are more specific concepts than “music”. People tag with different levels of specificity, and this usually reflects their level of expertise in the subject [9].

Extractions of POIs has so far been limited to users’ online tagging activities. However, many of the participants have authored papers which can be used to determine their research interests, and some of these interests are already available on `semanticweb.org` in the form of paper keywords. Acquiring such interests can be added to the system and used to improve recommendations on talks or sessions to attend, or people to meet. Also, information from social networking accounts can be used to avoid recommending existing friends.

Many users expressed their interest in retrieving their data after the conferences. The next version of LSS will give users permanent access to their LSS accounts, to enable them to revisit their logs of face-to-face contacts, to modify or regenerate their POIs, and to access all the services LSS provides. This will not only enable them to access their activity log, but it will also allow them to carry their accounts across conferences where this application is deployed.

More services will be provided in future LSS deployments, such as a ‘search for person’, ‘I want to meet’, and ‘find people with similar interests’. Data from RFIDs can be used to identify ‘best attended session or talk’. Social contacts from social networking systems and COPs could be used to find out who has made new contacts, especially if we can compare data over several LSS deployments.

8 Conclusions

The Live Social Semantics application is pioneering the full integration of active RFIDs with semantics and social networking systems. The paper described and evaluated the generation of Profiles of Interests for individuals by analysing their public tagging activities on Flickr, Delicious, and lastFm. The paper reported results from deploying this application at two international conferences, ESWC 2009 and HyperText 2009, during which 300 people took part in LSS. 236 users shared 371 SNS accounts on the LSS site. A POI was generated for each of these users, and saved by 72 of them. Overall, 21% of the interests suggested by our system were removed by users. When analysing logs of 11 HT09 users who clearly edited their POIs, we found that 15% of the interests they rejected were due to incorrect semantic association. Further research is required to better understand users’ rationale for removing/keeping interests, and for using their collective intelligence to improve our POI generation processes.

Acknowledgement

This research was supported by the TAGora project (Future and Emerging Technologies program of the European Commission, contract IST-34721) and WeGov (EU ICT-248512).

References

1. Alani, H., Szomszor, M., Cattuto, C., den Broeck, W.V., Correndo, G., Barrat, A.: Live social semantics. In: Bernstein, A., Karger, D.R., Heath, T., Feigenbaum, L., Maynard, D., Motta, E., Thirunarayan, K. (eds.) ISWC 2009. LNCS, vol. 5823, pp. 698–714. Springer, Heidelberg (2009)
2. Angeletou, S., Sabou, M., Motta, E.: Semantically enriching folksonomies with flor. In: Workshop on Collective Intelligence & the Semantic Web (CISWeb), ESWC, Tenerife, Spain (2008)
3. Barrat, A., Cattuto, C., Colizza, V., Pinton, J.-F., den Broeck, W.V., Vespignani, A.: High resolution dynamical mapping of social interactions with active RFID (2008), <http://arxiv.org/abs/0811.4170>
4. Begelman, G., Keller, P., Smadja, F.: Automated tag clustering: Improving search and exploration in the tag space. In: Proc. 17th Int. World Wide Web Conf., Edinburgh, UK (2006)

5. den Broeck, W.V., Cattuto, C., Barrat, A., Szomszor, M., Correndo, G., Alani, H.: The live social semantics application: a platform for integrating face-to-face proximity with on-line social networking. In: Workshop on Communication, Collaboration and Social Networking in Pervasive Computing Environments (PerCol 2010), IEEE Int. Conf. on Pervasive Computing and Communications (PerCom), Mannheim, Germany (2010)
6. Eagle, N., Pentland (Sandy), A.: Reality mining: sensing complex social systems. *Personal Ubiquitous Comput.* 10(4), 255–268 (2006)
7. Garcá-Silva, A., Szomszor, M., Alani, H., Corcho, O.: Preliminary results in tag disambiguation using dbpedia. In: Knowledge Capture (K-Cap 2009) - Workshop on Collective Knowledge Capturing and Representation - CKCaR 2009, CA, USA (2009)
8. Glaser, H., Millard, I., Jaffri, A.: Rkbexplorer.com: A knowledge driven infrastructure for linked data providers. In: Bechhofer, S., Hauswirth, M., Hoffmann, J., Koubarakis, M. (eds.) *ESWC 2008*. LNCS, vol. 5021, pp. 797–801. Springer, Heidelberg (2008)
9. Golder, S.A., Huberman, B.A.: Usage patterns of collaborative tagging systems. *Journal of Information Science* 32, 198–208 (2006)
10. Guy, M., Tonkin, E.: Tidying up tags? *D-Lib Magazine* 12(1) (2006)
11. Hayes, C., Avesani, P., Veeramachaneni, S.: An analysis of the use of tags in a log recommender system. In: *Int. Joint Conf. Artificial Intelligence (IJCAI)*, Hyderabad, India (2007)
12. Hui, P., Chaintreau, A., Scott, J., Gass, R., Crowcroft, J., Diot, C.: Pocket switched networks and human mobility in conference environments. In: *WDTN 2005: Proc. 2005 ACM SIGCOMM workshop on Delay-tolerant networking*. ACM, New York (2005)
13. Li, X., Guo, L., Zhao, Y.E.: Tag-based social interest discovery. In: *Proc. 19th Int. World Wide Web Conf. (WWW)*, Beijing, China (2008)
14. Mathes, A.: Folksonomies - cooperative classification and communication through shared metadata. In: *Computer Mediated Communication - LIS590CMC* (December 2004)
15. Passant, A., Laublet, P.: Meaning of a tag: A collaborative approach to bridge the gap between tagging and linked data. In: *Workshop on Linked Data on the Web (LDOW)*, Int. World Wide Web Conference, Beijing, China (2008)
16. Scherrer, A., Borgnat, P., Fleury, E., Guillaume, J.-L., Robardet, C.: Description and simulation of dynamic mobility networks. *Comput. Netw.* 52(15), 2842–2858 (2008)
17. Specia, L., Motta, E.: Integrating folksonomies with the semantic web. In: Franconi, E., Kifer, M., May, W. (eds.) *ESWC 2007*. LNCS, vol. 4519, pp. 624–639. Springer, Heidelberg (2007)
18. Szomszor, M., Alani, H., Cantador, I., O'Hara, K., Shadbolt, N.: Semantic modelling of user interests based on cross-folksonomy analysis. In: Sheth, A.P., Staab, S., Dean, M., Paolucci, M., Maynard, D., Finin, T., Thirunarayan, K. (eds.) *ISWC 2008*. LNCS, vol. 5318, pp. 632–648. Springer, Heidelberg (2008)
19. Szomszor, M., Cantador, I., Alani, H.: Correlating user profiles from multiple folksonomies. In: *Proc. Int. Conf. Hypertext (HT 2008)*, Pittsburgh, PA, USA (2008)
20. Tesconi, M., Ronzano, F., Marchetti, A., Minutoli, S.: Semantify del.icio.us: automatically turn your tags into senses. In: *Social Data on the Web, Workshop at the 7th ISWC* (2008)
21. Thibodeau, P.: IBM uses RFID to track conference attendees (2007), <http://pcworld.about.com/od/businesscenter/IBM-uses-RFID-to-track-confere.htm>
22. Wu, L., Waber, B., Aral, S., Brynjolfsson, E., Pentland, S.: Mining face-to-face interaction networks using sociometric badges: Evidence predicting productivity in it configuration. In: *The 2008 Winter Conference on Business Intelligence*, University of Utah (2008)
23. Yeung, C.-M.A., Gibbins, N., Shadbolt, N.: Tag meaning disambiguation through analysis of tripartite structure of folksonomies. In: *2007 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology Workshops*, pp. 3–6 (2007)