

LODifying personal content sharing

Oscar Rodríguez Rocha ^a, Fabio Mondin ^b, Carmen Criminisi ^c and Larent-Walter Goix ^c

^a *Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129. Turin, Italy.*

E-mail: oscar.rodriguezrocha@polito.it

^b *Matrix S.p.a, Piazza Einaudi 8, 20124. Milan, Italy.*

E-mail: fabio.mondin@matrix.it

^c *Telecom Italia S.p.a., Piazza Einaudi 8, 20124. Milan, Italy.*

E-mail: {carmen.criminisi, laurentwalter.goix}@telecomitalia.it

Abstract. The advent of contemporary mobile devices and their increasing computing power and location capabilities combined with the most innovative web technologies has provided mobile users with new possibilities to share experiences on-the-go. The growing quantity of multimedia content present on the web makes difficult for mobile users to retrieve suitable content. Typically, users looking for interesting content related to their current position or POI, access web engines relying on keywords to describe their ideas. Unfortunately such descriptions are often subjective and retrieval can be ineffective. To address these issues, our platform provides users with an application targeted for modern mobile devices that allows content acquisition and publication. Published content is automatically analyzed and stored on our server with semantic annotations based on users' context and content, for further semantic search. We describe how and why we migrated from a triple-tags technology to semantics, hoping for related Linked Data.

Keywords: Mobile User-generated Content, Linked Data, Semantic Web

Introduction

The increasing popularity of mobile devices and their continuously growing computing power and location capabilities together with the most innovative web technologies and services, offers web users new possibilities to share experiences on-the-go by allowing them to create multimedia content that can be easily distributed and published over the Internet.

On the other hand, as technological improvements are giving a boost to multimedia content creation process, finding suitable content is becoming more problematical for mobile users. Typically, mobile users looking for interesting content related to their current position or point-of-interest (POI), access web engines from their mobile devices relying on keywords to describe their intent. Unfortunately such descriptions are often subjective and retrieval can be ineffective. In fact, there is no point in making available multimedia information that can only be found by chance [1].

Enabling multimedia content retrieval based on semantics, increases the effectiveness and quality of the

results as semantically annotating a content or a piece of information gives the possibility to quickly access and use the immense knowledge offered by the web (user-generated or not).

This work describes a solution to address the previous issue. In the first section we provide a general overview of our platform. It consists of an application targeted for modern mobile devices to allows content acquisition and publication that is provided to users. Then it analyzes automatically the published content and stores it on our server enriched with tags based on user's context. For accessing content users rely on keywords to describe it.

Section 2 shows the work done to improve content retrieval features through the introduction of semantics. Now the content published by the user is automatically analyzed and stored on our server, enriched also with semantic annotations based on users' context and preferences. Specific interfaces for mobile and web are provided to allow semantic searching and browsing of

stored user-generated content and their related linked data.

Some works related to this research topic are listed in section 5. Finally, we present some insights of our future work and conclusions regarding our platform improvements.

1. Problem Statement

Our starting point is a platform for sharing mobile User Generated Content, based on the core of the Coppermine Photo Gallery¹. On top of it, we developed social features, improved content navigation and created a set of APIs for mobile publishing and retrieval.

Content that is uploaded to the system can be cross-posted to different popular sites and social networks (like Facebook², Flickr³ and Twitter⁴), tagged with keywords and in the case of pictures, it is also possible to create a graphical annotation over a particular section. Users can sign-in and avoid registration using their OpenID⁵ accounts of any OpenID provider.

Even if this platform offers innovative features, users are starting to perceive it as something belonging to the past. The people is getting used to personal content sharing with the inevitable consequence that it gets more and more difficult for them to retrieve it as easy as the would like, so that social networks are getting smaller, thematic and personal. Moreover the new challenge on social applications, is to understand the users' needs and give them what they want before they start searching for it. Semantic technologies can be the ace in the hole for content enrichment and retrieval in the social and UGC world.

1.1. Current Approach to Content Annotation

When users take a picture or create a video with their smartphones, our client application prompts for custom tags and for a title. The content, user-defined tags, title, associated timestamp and GPS coordinates (when available) are uploaded to our platform. To overcome problems of limited connectivity and battery management, the client supports a deferred content up-

loading procedure. Pictures, videos and related metadata are associated to their creation timestamp.

After being uploaded, each content is processed by the platform, which adds the user's context tags. Context data is retrieved from the context management platform that provides the location, nearby buddies and calendar entries associated to the moment in which the picture was taken. Focus is given on location: our platform converts GPS coordinates whenever available from the device, into civil addresses. Moreover, it handles the retrieval of user-defined location labels. Context tags are generated according to a triple tags specification to carry a semantic meaning. We suggested brand new namespaces (address, people) that we use in addition to geo-namespaces widely used on social sites, e.g. Flickr⁶.

Our platform stores published content with its associated metadata and provides also context-based content navigation. Tagged pictures and videos are organized in *virtual albums* generated dynamically. These tag-based collections exploit triple tags to organize content: it is therefore possible to filter user-generated pictures by each triple tag namespace, predicate or value. For instance, it is possible to show only the pictures taken by a specific user (e.g. `people:fn=Walter+Goix`), in a specific GSM cell (e.g. `cell:cgi=460-0-9522-3661`) or acquired in a crowded area (`place:is=crowded`).

Moreover, content can be syndicated as context-filtered feeds in order to enable social services. Triple tags "semantics" is used to optimize the platform GUI: context tags are displayed in a friendly format, and are separated from user-defined tags.

1.2. Opportunity Through Semantics

In a mobile scenario, searching and accessing content has to be as performing as possible. The added value for mobile users is not only to access content itself but also to retrieve the "best content" in a specific situation. If the target is a very effective retrieval process the focus must be, first, on the way to catalog content. If we think to Flickr, maybe the most known portal for user-generated content, we immediately notice that image classification is performed by users themselves, thus creating folksonomies.

Keyword-based searches, especially when relying on user-generated tags with wild-free vocabulary, re-

¹<http://coppermine-gallery.net/>

²<http://www.facebook.com>

³<http://www.flickr.com>

⁴<http://www.twitter.com>

⁵<http://openid.net/>

⁶<http://www.flickr.com>

strict the amount of retrievable content and the type of queries that can be performed. However, the main problem of such approach is the ambiguity in which one can incur: the thoughts of a tag creator in a specific situation can be very different of a tag consumer in the same situation. User classification through self-defined tags and tag-based search must trigger the “process of mutual understanding” that involves tag creators and consumers[2].

To try to overcome this ambiguity when attaching user-generated tags, an extra step is needed: annotate content with semantically meaningful labels referring to an ontology. The creation of domain ontology is fundamental in the semantic web context and can be greatly helpful (for example, in a context like tourism, where places, events and people are popular types of entities).

Every uploaded content can hence be stored with these kinds of new and richer tags that allow a better management and, accordingly, an effective access and retrieval of multimedia data.

The annotation mechanism can be distinguished between manual and automatic. Even though the latter is challenging. In both cases the process must be assisted by a predefined ontology in order to use a common way of annotation in the phases of cataloging and retrieval.

Besides helping in classifying data and content for search, annotation and semantic offer a powerful query mean that goes beyond tags and keywords. They help to deepen knowledge by allowing creation of structured linked data that can be mashed-up together in several ways based on the query situation.

2. Semantic Enhancement

2.1. Semanticizing the relational

Semantic technologies were applied to the relational database standing behind our system in order to rebuild and semanticize the data structure, so that contents can be retrieved by means of SPARQL queries (improving system search capabilities) and most of them, can be linked to data from other providers, allowing automatic data enrichment.

A deep analysis of the data structure of the platform was performed to select the most useful information in the database, avoiding service tables and focusing on the ones that describe content, users and their relationships. Once the information was identified, two

kinds of data to represent emerged: intra-table information and cross-table information. The semanticization process started from intra-table data: in a relational database, every table has a primary key field, which is unique by definition, so it can be used for constructing the URI of the resource described by this table. For each resource, the information is stored in the other columns of the table, so it was necessary to find an appropriate predicate to construct a triple. Once all the intra-table information was processed, every resource got its own URI and once again, interlinking was just a matter of finding the appropriate predicate.

This URI and triple construction procedure (once the tables, columns and predicates to use have been chosen), can be easily made by means of the D2R server⁷. Usually, the D2R server runs as a service that enables SPARQL queries by working "on top" of a relational DB, but in our case, we used its *dump-rdf feature* to write a mapping file in n-triples format which once completed, allows the creation of a semantic database dump in n-triple format following the mapping described in the mapping file. The reason standing behind this choice, resides in the final result we want to obtain: a semantic platform offering Linked Data functionalities and running locally. All the triples coming from our platform, together with external data coming from the principal data providers (DBpedia⁸, Geonames⁹ and Linkedgeodata¹⁰) was then inserted into our semantic triple store server (Openlink Virtuoso¹¹) which in its commercial edition offers advanced querying functions (such as geo-localized sparql query and text-search) that can be very useful in context of social applications.

2.1.1. Keywords and Semantic Annotator, a link to Linked Data

Keywords were one of the critical aspects of the semanticization process. In our system, there was no need to have access to each single keyword at a database level, so all the keywords of a resource were saved in a single column (space-separated). In a semantic application an "all keywords" information is not useful, so we had to separate all keywords and make triples describing each one. Keywords are important because they can be used as an hook to link external resources to our data: if for example, a content is tagged with the

⁷<http://www4.wiwiss.fu-berlin.de/bizer/d2r-server/>

⁸<http://dbpedia.org>

⁹<http://geonames.org>

¹⁰<http://linkedgeodata.org>

¹¹<http://virtuoso.openlinksw.com>

word "Coliseum", this keyword can be used to retrieve information about "The Roman Colosseum" in the external datasets and use triples to say that the content refers to it. This idea opens some issues about semantic annotation: the association between a keyword and a concept is not always so easy as it is for the "Coliseum" case, usually it is necessary to perform disambiguation. Even if in some way this is still an open problem for which we are evaluating different solutions (such as Apache Stanbol¹²), we developed a semantic annotator for this purpose (based on SMOB¹³), which performs a basic named entity resolution, language detection and concept retrieval with ranking, we use it to perform the Linked Data linking.

2.2. Automatic Semantic Tagging/Analysis

On the way to semantically cataloging the mobile user-generated content sent by users, the challenging goal of this task is to contextualize the content by analyzing any related information. As previously described, some information is input by the user himself (e.g. title and/or tags) whilst other is automatically derived by our context management platform, such as user location or nearby friends. This step aims at associating a unique semantic LOD (Linked Open Data) concept (and related resource) to every possible set of information. Whilst it may seem easier to link places and people to specific non-ambiguous resources, many challenges arise when trying to identify a "best resource" that maps to a single user-generated tag or word within a title. Such challenges are still common issues with text mining and natural language processing techniques and this paper does not intend to provide a systematic solution to this problem, rather a reasonably fair approach to the main issue of semantically cataloging user-generated content. This solution is described below in details.

2.2.1. Location Analysis

Whenever a content is received on our platform, its sender is identified and, upon success, possibly contextualized. The provided output, whenever available, results in a location (GPS coordinates, civic location and user-labeled place name and type when applicable), and a list of nearby friends (user names and full names). The location information also contains a refer-

ence to a Geonames¹⁴ resource, which validity is guaranteed by the locationing process itself based on the geocoordinates. In particular, the (nearest) city-level resource is returned.

Each nearby friend is referenced by a local descriptive RDF resource that also links to external identifiers that the user declared on our platform (such as the user accounts on popular Social Networks). We evaluated the linking of nearby users to existing "external" resources based on the results of a Sindice¹⁵ query: however as the results may be ambiguous and may trigger privacy concerns this further automatic process was turned off and only local linking was retained.

The mobile application further provides a way for the user to associate a content to a specific POI (e.g. monument). This feature is based on a search provider on our platform, which returns local POIs out of a list of information providers (e.g. Google Local¹⁶) based on the identified location of the user. When the user explicitly associates a content to a POI, a specific tag is sent to the platform with a reference to that POI (e.g. poi:recs_id=72). This module analyzes the tag and based on the POI name, category and location derived from the platform, tries to identify the related DBpedia resource using SPARQL. At this time commercial categories such as restaurants, hotels, etc are excluded from this analysis. Below is an example of SPARQL query to achieve this result:

2.2.2. Text Analysis

Title and plain tags analysis follow a well-more complex process to provide candidate resources.




Figure 1. Semantic annotation process.

In particular, semantic resource extraction from title analysis is illustrated in Figure 1 and performs as follows: First, a text processing is performed. The title language is initially identified using [3] based on [4]. Then, a morphological analysis is performed using

¹²<http://incubator.apache.org/stanbol/>

¹³<http://www.smob.me>

¹⁴<http://www.geonames.org>

¹⁵<http://www.sindice.com>

¹⁶<http://www.google.com/lochp>

FreeLing[5] configured with the identified language. FreeLing was preferred over TreeTagger[6] as it allows for multiwords lemmas detection and potentially provides WordNet sense annotation and Named Entity classification, although not yet on all supported languages (whilst it performs morphological analysis on all languages of our use case). From that analysis, NP (Proper Nouns) lemmas are extracted whilst other part-of-speech are discarded. This filtering process can be related to an eTourism use case, which typically gains interest in identifying proper nouns (such as places, sights, events, people). At this time, non-numeric NP lemmas with a score of at least 0.2 are preserved and merged with plain tags to compute a well-defined list of unique (multi)words. We do understand that nouns or verbs can be useful to describe a peculiar characteristic of the content or the place it was taken (e.g. the subject of a picture or the activity that can be done there) although a further pruning would be required to restrict to concrete concepts only, further discarding abstract statements (e.g. “difference”, “joyness”, etc). At this stage, we thus use term frequency to further process the title and extract other potential relevant words. However, we intend to use the WordNet sense annotation capability of FreeLing for this purpose in the future, coupled with disambiguation techniques and others to select only lemmas.

The next step involves a semantic brokering component. This component is assisted by a set of resolvers that perform full-text or term-based analysis based on the previous output. Such resolvers are aimed at providing candidate semantic concepts referring to Linked Open Data as well as additional related information if available. Resolvers may be domain- or language-specific, or general purpose. For term-based analysis, each word of the previously-computed list is individually processed to identify a list of candidate LOD resources to match with. Such activity is accomplished by leveraging techniques provided by SMOB¹⁷. In particular, we rely on the provided framework for invoking a set of predefined services, such as DBpedia¹⁸ and Sindice, further extended to Evri¹⁹ entity resolver. In particular, whilst most of such resolvers provide web service type of interfaces, DBpedia query has been optimized to rely on SPARQL rather than of its lookup service. This functionality allows to benefit from the full-text support, as well as additional filters e.g. based on

language, entity type & native scoring. The query also follows resource redirections to avoid returning “disambiguation” pages. In addition, we further realized that in some cases Named Entity Recognition would benefit from the original context (the whole title) to help disambiguation. As such we also rely on full-text based resolvers such as Evri and Zemanta²⁰ to derive additional candidates.

Next, a semantic filtering applies to process candidate LOD resources received by the broker and perform disambiguation. Resolvers return many candidate resources that eventually match with the actual concept behind the single word. Whilst the disambiguation task is humanly solved in the case of semantic search and browsing of content where a dynamic user interface is proposed to the user for selection, our goal is to automatically select and discriminate the most appropriate candidate resource. This disambiguation process starts with a priority and validation mechanism to filter and rank the candidates resources extracted. Such mechanism is based on the analysis of the target resources retrieved and their associated ontologies/graphs. We decided that resources referring to Geonames graph have higher priority than the ones related to DBpedia, followed by Evri types of resources. At this time all candidate resources pointing to other graphs are discarded. Note that for some resolvers, e.g. Sindice, candidate resources may refer to various ontologies, e.g. Geonames or DBpedia or others. For this reason we associate priorities with graphs and not with the resolvers.

The reason we give Geonames the highest priority is mainly for consistency with our contextualization process in which locations are systematically linked to a Geonames resource. Second, besides being very exhaustive on locations, Geonames is also focused on this types of resources where very little overlap exists with other types of resources. For this reason we prefer it over DBpedia, which remains our preferred graph for generic and heterogeneous concepts, especially in the context of tourism, points-of-interest and celebrities, where rich and well-linked data is available.

Besides this priority mechanism, a validation is performed to check whether the resource itself is valid. This step depends on the single ontology, which provides hints for validating the resource. For example, for DBpedia resources, we query the SPARQL²¹ endpoint to check whether they contain an actual binding,

¹⁷<http://www.smob.me>

¹⁸<http://www.dbpedia.org>

¹⁹<http://www.evri.com>

²⁰<http://www.zemanta.com/>

²¹<http://www.w3.org/TR/rdf-sparql-query/>

or whether they correspond to a disambiguation page whenever the “disambiguates” property exists (only when the candidate resource is not coming from DBpedia resolver, which already performs a similar disambiguation check). Resources that do not validate are discarded. We plan to add semantic checks (e.g. properties of the resource, cross-resource type and relation consistency) as well as WordNet analysis for nouns in the future.

For automatic annotation, a final check is performed by calculating string similarity between the remaining list of candidates and the original corresponding words or lemmas. In particular, after initial empirical tests, candidates with Jaro-Winkler distance lower than 0.8 are discarded at this stage unless their DBpedia score is maximum. Automatic annotation is performed only in case a single candidate remains after this step, to avoid ambiguity and limit errors. Empirical tests proof that such technique must be further improved as it still provides false positives.

2.3. Semantic Virtual Album

We propose a new semantic-based navigation approach for content retrieval, which provides users with the possibility to reach the best content by exploiting their current situation and context information. Give users the best content in a specific context, means that we have both to provide a strong search instrument and anticipate the users’ retrieval needs and that is what our virtual albums do.

A virtual album is a collection of multimedia objects retrieved dynamically by applying several complex search conditions over our data storage. Organizing dynamically multimedia content into virtual albums based on keywords or by simple criteria such as “by user”, “by date” or “by place” was something already possible by means of relational DB technology, but with the use of semantic technologies both search (by allowing the construction of complex search criteria with SPARQL²²) and data enrichment (by exploiting the Linked Open Data) processes can be significantly improved.

Once the automatic semantic analysis is performed (as described in the previous paragraph), the uploaded content is saved in our Semantic Triple Store (An Openlink Virtuoso Commercial Edition installation running on a 32bit dual core processor Virtual ma-

chine with 12GB RAM), together with other Linked Open Data imported datasets (such as DBpedia and Geonames). The creation of a virtual album is performed through SPARQL queries made on our Triple Store endpoint and by managing the content rendering: SPARQL is used to express queries across several datasets and its expressiveness helps creating “complex” queries that are not allowed by the traditional keyword search. Such queries can be richer, more elaborated and accurate as several matching clauses can be specified to retrieve specific resources and related information, also relying on inference capabilities.

In order to leave the possibility to our users to retrieve a set of content defined as:

Select the set of user generated content, taken near to the monument “Mole Antonelliana”

1. PREFIX rdfs:http://www.w3.org/2000/01/rdf-schema#
2. PREFIX sioc:http://rdfs.org/sioc/types#
3. PREFIX comm:http://comm.semanticweb.org/core.owl#
4. PREFIX rev:http://purl.org/stuff/rev#
5. SELECT DISTINCT ?link WHERE {
6. ?monument rdfs:label "Mole Antonelliana"@it .
7. ?monument geo:geometry ?sourceGEO .
8. ?resource geo:geometry ?location .
9. ?resource a sioc:MicroblogPost .
10. ?resource comm:image-data ?link .
11. FILTER(bif:st_intersects(?location, ?sourceGEO, 0.3)) .
12. }

To understand how this works, let us start from line 6, in which the query selects a monument resource labeled “Mole antonelliana” (in case of our triple store, the related DBpedia resource is returned). On line 7, we take the location with the predicate geo:geometry, which is a specific property that has been previously associated by means of Virtuoso platform capabilities²³. A geometry is a position specification which says that a particular object stands over a particular area.

On line 8 we start querying the user generated data contained in our store. Line 8 selects all geometry-referenced entries, line 9 filters them by selecting only the ones user-generated and line 10 selects from the RDF descriptor the URL of the real content.

²²<http://www.w3.org/TR/rdf-sparql-query/>

²³More info can be found on:
http://docs.openlinksw.com/virtuoso/fn_rdf_geo_fill.html and
<http://docs.openlinksw.com/virtuoso/rdfsparqlgeospat.html>

Line 11 is a call to a specific Virtuoso geo-related function, which filters the data by distance from the “Mole Antonelliana” monument.

This query can become more complex, for example by filtering the content taking into account the relationships between users (social filtering):

Select the set of user generated content, taken near to the monument “Mole Antonelliana”, created by users who are friends of user “oscar”

1. PREFIX rdfs:
http://www.w3.org/2000/01/rdf-schema#
2. PREFIX sioc:http://rdfs.org/sioc/types#
3. PREFIX comm:
http://comm.semanticweb.org/core.owl#
4. PREFIX rev:http://purl.org/stuff/rev#
5. PREFIX foaf:http://xmlns.com/foaf/0.1/
6. SELECT DISTINCT ?link WHERE
7. {
8. ?monument rdfs:label "Mole Antonelliana"@it .
9. ?monument geo:geometry ?sourceGEO .
10. ?resource geo:geometry ?location .
11. ?resource a sioc:MicroblogPost .
12. ?resource comm:image-data ?link .
13. ?resource foaf:maker ?user .
14. ?oscar foaf:name "oscar" .
15. ?user foaf:knows ?oscar .
16. FILTER(bif:st_intersects(?location, ?sourceGEO, 0.3)) .
17. }

On line 13 we select the author of the selected UGC, whilst on line 14 and 15 we filter the user who knows user “oscar”.

The query can be refined once more by using the rating information to order or filter the content, thus further restricting to highly-rated content:

Select the set of user generated content highly rated, taken near to the monument “Mole Antonelliana”, created by users who are friends of user “oscar”

1. PREFIX rdfs:
http://www.w3.org/2000/01/rdf-schema#
2. PREFIX sioc:http://rdfs.org/sioc/types#
3. PREFIX comm:
http://comm.semanticweb.org/core.owl#
4. PREFIX rev:http://purl.org/stuff/rev#
5. PREFIX foaf:http://xmlns.com/foaf/0.1/
6. SELECT DISTINCT ?link WHERE {
7. ?monument rdfs:label "Mole Antonelliana"@it .

8. ?monument geo:geometry ?sourceGEO .
9. ?resource geo:geometry ?location .
10. ?resource a sioc:MicroblogPost .
11. ?resource comm:image-data ?link .
12. ?resource foaf:maker ?user .
13. ?oscar foaf:name "oscar" .
14. ?user foaf:knows ?oscar .
15. ?resource rev:rating ?points.
16. FILTER(bif:st_intersects(?location, ?sourceGEO, 0.3)) .
17. }
18. ORDER BY DESC(?points)

As anticipated above, behind a virtual album stands a SPARQL query, which is able to retrieve the searched content dynamically with very precise search criteria. However, it is also clear that to exploit at best the possibilities of semantic technologies we must be able to present the virtual albums to the final user in the right way, especially in mobility.

3. Web interface

The platform’s web interface offers users an environment to perform many operations: from personal profile and social features management to content browsing or advanced content editing. It’s targeted for modern web browsers and when it is accessed from a mobile device, redirects the user automatically to the mobile interface (giving also the possibility to switch back to the normal web interface).

4. Mobile interface

Our mobile web interface aims to be simple and easy to use for mobile users. When the latter is accessed, it will first try to retrieve the user’s geographical position (a user authorization dialog will be prompted) by using the mobile browser’s location API. With this information, the possibility of filtering geographically the results can be provided to the user. Then, the interface is presented with a search box, where users can input keywords for retrieving a desired content.



Figure 2. Search box and user's location coordinates.

The search field is automatic and AJAX-based²⁴, which means that each time, 2 seconds after the last keystroke is pressed, a query is performed and a list of candidate results will be displayed.

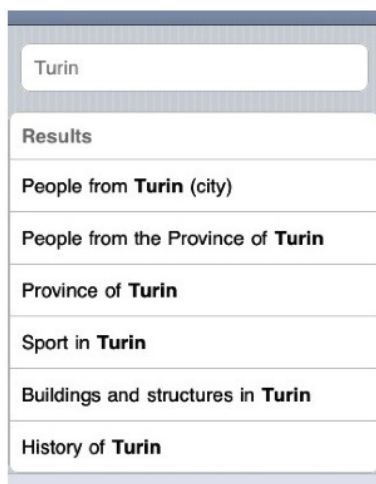


Figure 3. Result candidates are listed for “Turin”.

The user can click on the result that matches his search to visualize all the content associated with the selected resource that will be presented in a list view with thumbnail previews of the content, brief descriptions and clickable links for a more detailed view of each single element.

This view further shows an “About” button on top, which can display a mashup of information related the selected resource based on the LOD datasets.



Figure 4. Content associated to the selected resource and the “About” button.

4.1. Enrichment and mashup

As previously said, the semantic query mechanism allows combining data coming from different sources. Data presentation can hence be further enhanced and enriched with linked data coming from the semantic web, and is achieved here by clicking the “About” button.

Our main information source is DBpedia, where rich, structured and well-linked data is available. Indeed queries to this source allow collecting interesting and meaningful linked data that is then filtered based on which information and/or content has to be shown to the user. Here is an example of an SPARQL query that enables the mashup creation:

1. PREFIX dbpo: <http://dbpedia.org/ontology/>
2. PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#>
3. PREFIX lgdo: <http://linkedgeodata.org/ontology>
4. PREFIX tl-pid: http://beta.teamlife.it/cpg148_pictures
5. PREFIX comm: <http://comm.semanticweb.org/core.owl#>
6. SELECT DISTINCT ?lbl ?entType ?desc ?others
7. WHERE { {
8. tl-pid: ".\$this->picture_details->pid." geo:geometry ?locPID .
9. ?city geo:geometry ?locCity .
10. ?city a ?entType .
11. ?city rdfs:label ?lbl .
12. ?others dbpo:abstract ?desc .

²⁴[http://en.wikipedia.org/wiki/Ajax_\(programming\)](http://en.wikipedia.org/wiki/Ajax_(programming))


```

13. ?others a dbpo:Place .
14. FILTER (?entType in (lgdo:City)) .
15. FILTER langMatches(lang(?desc), 'it').
16. FILTER( bif:st_intersects( ?locPID, ?locCity, 1 )
).
17. } }LIMIT 5
18. } UNION
19. {SELECT DISTINCT ?lbl ?entType ?desc ?oth-
ers WHERE
20. { {
21. tl-pid:" . $this->picture_details->pid . " geo:geometry
?locPID .
22. ?others geo:geometry ?location .
23. ?others a ?entType .
24. ?others rdfs:label ?lbl .
25. OPTIONAL{
?others <http://linkedgeodata.org/property/website>
?desc} .
26. FILTER (?entType in (lgdo:Restaurant ) ) .
27. FILTER( bif:st_intersects( ?locPID, ?location,
0.3 ) ) .
28. } }LIMIT 5 }
29. UNION
30. {SELECT DISTINCT ?lbl ?entType ?desc ?oth-
ers WHERE
31. { {
32. tl-pid:" . $this->picture_details->pid . " geo:geometry
?locPID .
33. ?others geo:geometry ?location .
34. ?others a ?entType .
35. ?others rdfs:label ?lbl .
36. OPTIONAL {
?others <http://linkedgeodata.org/property/website>
?desc} .
37. FILTER (?entType in (lgdo:Tourism)) .
38. FILTER( bif:st_intersects( ?locPID, ?location, 1
))
39. } }LIMIT 5 } UNION
40. {SELECT DISTINCT ?lbl ?entType ?desc ?oth-
ers WHERE
41. { {
42. tl-pid:" . $this->picture_details->pid . " geo:geometry
?locPID .
43. ?others geo:geometry ?location .
44. ?others a ?entType .
45. ?others rdfs:label ?lbl .
46. ?others comm:image-data ?desc.
47. FILTER (?entType in (sioc:MicroblogPost)) .
48. FILTER( bif:st_intersects( ?locPID, ?location,
0.2 ) )
49. } }LIMIT 5 } }

```

With this query, starting from a picture sent to our system by the tourist and its semantic location information, useful information is retrieved for the user such as the description (from DBpedia) of the city where the tourist is, the restaurants (and their websites) near the user's location and other touristic attractions in the vicinity that the mobile users can decide to visit, and other UGC content taken in the same location from other users.

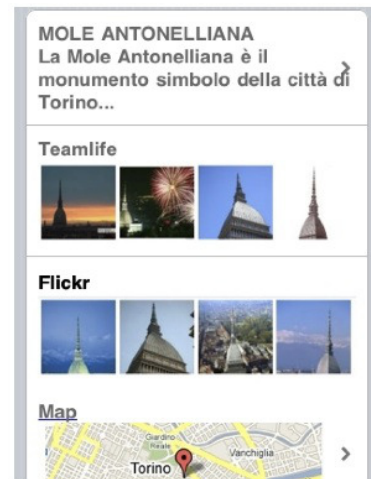


Figure 4. Mobile Interface mashup example.

5. Related Work

Many research activities have been done to provide solutions for information and user-generated content management in the eTourism field.

Garcia *et al*[7], has created a platform which allows the generation, annotation and retrieval of on-the-move multimedia tourist content while on a tourist destination. With a mobile application, a tourist can create content and annotate it based on multilingual folksonomies. An authorized person or body filters it before being published using a web page. This GUI allows browsing for content provided by either individual users or those including a tag or keyword associated to a concept. Generally this work addresses a similar scenario, but the semantic approach is focused only on supporting multilingual annotation through unique concepts, without mentioning semantic meaning extraction or concept disambiguation.

Zhang *et al*[8], has created a method to demonstrate how Semantic Web can support information integration and facilitate the creation of semantic mashups by integrating the Semantic Web with Web 2.0 in the field

of travel planning. We believe the same vision can also be applied to user-generated content, thus, their ideas contributed to inspire our work.

In the area of blogging, Zemanta²⁵, a content suggestion engine, uses natural language processing and semantic search technology to suggest pictures, tags and links to related articles mainly from social sites. This work inspired us on how to analyze user-generated content automatically.

A very interesting work has been done by Passant *et al*[9], with SMOB, a microblogging system that provides ways to leverage microblogging according to the Linked Data Web guidelines, providing semantic annotation tools. This work, together with SparqlPuSH [10] has strongly influenced and inspired our platform semantic enhancements, in particular for the concepts of semantic annotation brokering and of SPARQL-based virtual albums.

6. Conclusion and future work

We presented an approach to introduce and implement semantic capabilities into our existing platform. Mobile users can benefit from automatic semantic enrichment of their content generated on-the-go, which also facilitates the content retrieval process.

As mentioned before, adding automatic analysis capabilities to annotate published content helps us to improve the performance of the content retrieval process and browsing. However, there's a huge amount of content already present in our platform that remains to be semantically annotated. Solving this issue requires to create and introduce new automatic batch processing mechanisms. As the user-assisted disambiguation is not used, it becomes more challenging to guarantee the right semantical meaning extraction, and user evaluations are planned to evaluate and improve our disambiguation algorithms.

Together with the increasing use of contemporary centralized social networks and web applications, the growth of user-generated content over the web, some issues regarding interoperability and decentralization, privacy and identity and trust building have emerged [11].

Given this scenario, we envision a federation of interconnected social networks and web applications, each one hosted right inside the end-users' home network

devices to provide users with full control of their content and information, customization, scalability and interoperability.

At the time of writing, we're planning to move our platform from a centralized to a federated architecture. This step will require so much effort and the need of mature technologies to make this change possible. However our target architecture is composed by different components:

6.1. Home network device

A generic NAS server attached to the user's home network will play a central role in the system: it will run the platform, store and stream users' content to different devices. It must run an embedded version of the Linux²⁶ operating system and some open software, protocols and server applications:

- A LAMP (Linux, Apache HTTP Server²⁷, MySQL database server²⁸ and PHP²⁹) web server hosts the federated social network application, making it available on the web through the use of a DDNS³⁰ mechanism that maps the home network dynamic IP address with a user's custom domain name.
- By implementing the UPnP³¹ set of networking protocols, the home network device can be set as UPnP media server to share the media content stored in the file system by the users' of the social network with any UPnP-compatible device.

6.2. Federated web application

The platform will have the possibility to be associated to a domain name and each member of the family can have their own account. The main features delivered should be:

- Profile data sharing and relationships with another networks, implemented with FOAF.
- A Webfinger³² protocol implementation enables the identification of users across different social networks and the identity validation.

²⁵<http://www.zemanta.com>

²⁶<http://www.linux.org/>

²⁷<http://httpd.apache.org/>

²⁸<http://www.mysql.com/>

²⁹<http://www.php.net/>

³⁰<http://tools.ietf.org/html/rfc2136>

³¹<http://www.upnp.org/>

³²<http://code.google.com/p/webfinger/>

- Multimedia content sharing, accomplished by using OEmbed³³.
- A users' activities timeline in the ActivityStreams³⁴ format.
- Publish and subscribe mechanism implemented through the PubSubHubBub³⁵ open protocol which also provides near-instant notifications.
- A Salmon³⁶ protocol implementation to comment and annotate the original sources of updates and content.
- Desktop and mobile web interfaces to browse the application.
- A set of OpenSocial³⁷ based API's to allow system extensibility.

6.3. Home devices

UPnP-compatible home devices can directly communicate with the home network device through the UPnP media server: they will be able to browse for available content on the media server and request a file for playback. For example, a UPnP-compatible photo-frame displaying a real-time slideshow of the media content that a family member is taking during his holidays, then when arriving home, they can watch together the pictures and videos on a UPnP-compatible TV.

³³<http://oembed.com/>

³⁴<http://activitystrea.ms/>

³⁵<http://code.google.com/p/pubsubhubbub/>

³⁶<http://www.salmon-protocol.org/>

³⁷<http://code.google.com/apis/opensocial/>

References

- [1] G. Stamou and S. Kollias, eds., *Multimedia Content and Semantic Web: Methods, Standards and Tools*. Chichester, West Sussex, Hoboken, NJ: John Wiley & Sons Inc., 2005.
- [2] Marinchev, "Practical semantic web – tagging and tag clouds," *Cybernetics and Information Technologies*, vol. 3, 2006. has not read yet.
- [3] "Lang detect." <http://pear.php.net/manual/en/package.text.text-languagedetect.php>.
- [4] W. B. Cavnar and J. M. Trenkle, "N-gram-based text categorization," in *Proceedings of SDAIR-94, 3rd Annual Symposium on Document Analysis and Information Retrieval*, (Las Vegas, US), pp. 161–175, 1994.
- [5] "Freeling." <http://www.lsi.upc.edu/nlp/freeling/>.
- [6] "Tree tagger." <http://www.ims.uni-stuttgart.de/projekte/corplex/TreeTagger/>.
- [7] A. Garcia, M. T. Linaza, I. Arbulu, I. Torre, and Y. Cobos, *A mobile platform for on-the-move user-generated tourist contents*, pp. 58–65. 2008.
- [8] L. Zhang, Q. Zhu, W. Wang, and P. Zhang, "Web 2.0 and Semantic Web Based Intelligent Service System for Tourism," *Wireless Communications, Networking and Mobile Computing, 2008. WiCOM '08. 4th International Conference on*, pp. 1–4, Sept. 2008.
- [9] A. Passant, T. Hastrup, U. Bojars, and J. Breslin, "Microblogging: A semantic and distributed approach," in *Proceedings of the 4th Workshop on Scripting for the Semantic Web, Tenerife, Spain, June 02, 2008, CEUR Workshop Proceedings*, 2008.
- [10] A. Passant and P. Mendes, "sparqlPuSH: Proactive notification of data updates in RDF stores using PubSubHubbub," vol. 699 of *CEUR Workshop Proceedings ISSN 1613-0073*, February 2010.
- [11] K. Falahi, Y. Atif, and S. Elnaffar, "Social networks: Challenges and new opportunities," in *Green Computing and Communications (GreenCom), 2010 IEEE/ACM Int'l Conference on Int'l Conference on Cyber, Physical and Social Computing (CPSCom)*, pp. 804–808, dec. 2010.