

Facilitating Dialogue - Using Semantic Web Technology for eParticipation

George Anadiotis¹, Panos Alexopoulos¹, Konstantinos Mpaslis¹,
Aristotelis Zosakis¹, Konstantinos Kafentzis¹, and Konstantinos Kotis^{2,3}

¹ IMC Technologies, Athens, Greece

{ganadiotis,palexopoulos,kmpaslis,azosakis,kkafentzis}@imc.com.gr

² University of the Aegean, AI-Lab, Samos, Greece

kotis@aegean.gr

³ Prefecture of Samos, IT Department, Samos, Greece

kotis@samos.gr

Abstract. In this paper we describe the application of various Semantic Web technologies and their combination with emerging Web 2.0 use patterns in the eParticipation domain and show how they are used in an operational system for the Regional Government of the Prefecture of Samos, Greece. We present parts of the system that are based on Semantic Web technology and how they are merged with a Web 2.0 philosophy and explain the benefits of this approach, as showcased by applications for annotating, searching, browsing and cross-referencing content in eParticipation communities.

1 Introduction: On eParticipation

There is general agreement among many actors in democratic societies that there is a lack of possibilities for political participation. New media and the internet especially have been heralded as solutions to this problem and while eParticipation, or 'the use of information and communication technologies to broaden and deepen political participation by enabling citizens to connect with one another and with their elected representatives' [12], cannot replace other forms of participation and solve the democratic deficit once and for all, it can help to improve the overall democratic culture in a society.

Even though the uptake of eParticipation solutions depends to a great extent on political commitment and social consensus, a solid technical foundation that provides advanced functionality via a seamless user experience and helps dealing with eParticipation-associated challenges [12] is also of great importance. With that goal in mind, when faced with the challenge to develop and deploy an eParticipation solution for the Regional Government of the Prefecture of Samos, we chose an approach powered by the advanced features of Semantic Web technology, while hiding its complexity behind a user-friendly Web 2.0 facade. The resulting solution has been deployed at <http://www.samos-dialogos.gr>.

This paper is organised as follows: in this Section, we present the domain and give the necessary background in order to explain the use of ontologies

in the deliberation process in Section 2, while in Section 3 we elaborate on the mechanisms in place for transparent hybrid-based search and browsing. In Section 4 we describe the reference deliberation ontology developed as part of the project and how it is used to make content available as Linked Data and build distributed applications that use it. Finally, we evaluate our approach in Section 5 and present conclusions and future work in Section 6.

1.1 A Methodology and Platform for eParticipation

In order to facilitate eParticipation, a two-fold approach has been taken, laying the theoretical framework and providing its implementation:

- a *methodology* has been developed, based on extensive study of related literature and deployed projects worldwide
- an *integrated platform* has been implemented, using appropriate tools to support the methodology in every aspect

The methodology is based on three different processes that can run separately, complementing each other and covering online channels of citizen participation. The first is the eConsultation process, initiated in a top-down way by the decision makers, aiming to provide an efficient and easy way to collect intelligent feedback from citizens on different issues. The second process is the ePetition process, initiated in a bottom-up way by citizens, aiming to offer an accessible mechanism through which citizens can raise publicly issues they deem important. Finally, the third process is the eDeliberation process, which requires a closer collaboration between decision makers and citizens in order to formulate policies and achieve consensus on important issues. All three processes have been implemented in the platform, supplemented by offline actions where needed.

In the rest of the paper we focus on eDeliberation, as it is the most engaging and complex process. It essentially constitutes a tight 'serial process' within a specific time-frame, with 7 concrete steps embedded in each deliberative cycle (see Figure 1). Initially we have the agenda setting stage, in which citizens may choose among a list of proposed issues for deliberation. After choice has been made, we move on to the discussion stage, succeeded by a report publishing stage, in which moderators summarize the discussion in a report. Then there is the voting stage, in which citizens have to fill in questionnaires to quantify opinions and finally, after another report has been published to summarize the outcome of the voting stage, an interactive online real-time council takes place. The process concludes with the publishing of an overall report.

The platform supports the implementation of this process via appropriate web-based tools for each step: an e-poll component for agenda setting, a moderated e-forum, e-surveys to quantify opinions, interactive online real-time web-casting of the council, where citizens can offer their feedback directly and a document library component for report publishing. Both the methodology and the platform have evolved through previous experience in the case of the City of Trikala, Greece [9], a finalist in the European eGovernment Awards for 2009.¹

¹ <http://www.epractice.eu/en/cases/edialogosawards>



Fig. 1. Methodology stages for eDeliberation

2 Adding Semantics to eDeliberation

Each deliberation process is focused on a specific domain, determined by the outcome of the agenda setting stage. In order to assist citizens engage in the process, as well as administrators responsible for producing information material for it, domain knowledge can be captured and made available to both user groups by means of a domain ontology. The association of a domain ontology with a deliberation is an optional task, however doing so serves a dual purpose: i) users are given context for the topic under discussion, in a form they can use to annotate content and ii) based on provided annotations and ontological relations, a powerful hybrid search mechanism is implemented.

2.1 Making Domain Knowledge Explicit: Developing Domain Ontologies

So far, four online deliberation processes have been started in the Prefecture of Samos, for which three ontologies have been developed, as two of the deliberations refer to the same domain (Tourism) and are hence associated with the same ontology. The ontologies have been developed by knowledge workers in the Prefecture of Samos, in collaboration with the University of Aegean, AI-Lab. The methodology followed for the development of the domain ontologies was partially based on HCOME [11]. Due to the short time of project implementation, the collaborative part of the methodology was not implemented using a tool such as SharedHCONE or Collaborative Protege, however active involvement of knowledge workers in the ontology engineering process was ensured.

Knowledge was extracted from administration documents provided by the Prefecture (press docs, board decisions docs, legislation, etc), web documents retrieved by knowledge workers and other available domain ontologies. Extraction of conceptualizations from text was performed manually, with all artefacts named in English while providing Greek labels for internationalization purposes. A small number of knowledge workers with less experience in ICT was asked to use a human-centered ontology engineering environment, namely HCONE2 [10], for the engineering of draft versions of the ontologies, while the final versions were developed in Protege.

The domain ontologies developed following the abovementioned methodology were:

- **Recycling Ontology**, aiming to represent knowledge regarding the domain of recycling materials and its impact to the environment. Special focus on the recycling of plastic bags from local companies/organizations has been given due to related administrative board decisions that were taken.
- **Culture Ontology**, aiming to represent knowledge regarding the cultural heritage of the prefecture and the role of the authorities and non-profit organizations to the accentuation of local assets, as well as potential initiatives towards new forms of cultural events/developments.
- **Tourism ontology**, aiming to represent knowledge regarding the touristic development of the prefecture and the role of the authorities and non-profit organizations to the accentuation of local assets as well as potential initiatives towards new forms of touristic developments.

2.2 Making Explicit Knowledge Accessible: Associating Domain Ontologies with the Deliberation Process

Even though knowledge formulation in the form of ontologies can provide substantial benefits to the operation of the platform, its potential will not be fully realized unless it is available to the users in a form that is familiar and easy to use. For this purpose, we chose to pursue a Web 2.0 approach as far as user experience is concerned.

Ontologies that the administrators wish to associate with certain deliberations are acquired (developed from scratch or imported/reused) and stored in an ontological repository. Ontologies are processed before being stored in the repository, in order to ensure that all classes and instances have corresponding labels that are readable and informative.

After a deliberation has been created in the system, administrators may use a GUI to browse the ontological repository and select an ontology they wish to associate with the deliberation. The ontology is then imported in the system in the form of a taxonomy: a stripped-down form of the original ontology, retaining only classes and instances as well as their hierarchical relations, using their labels as taxonomy terms. The ontology will be from that point on available across all the tools encompassed in the deliberation in the form of a taxonomy of domain-related terms that can be used to:

- **Annotate.** A deliberation's domain taxonomy can be used to annotate content associated with the deliberation (web content, forum posts, documents, etc). Users remain agnostic of the underlying complexity of the ontology and are presented with a tree-like structure with auto-complete assistance that they can use to 'tag' resources, as seen in Figure 2.
- **Browse.** In order to enable quick access to popular content within a deliberation, users may browse through a tag cloud used to visualize resource annotation frequency. Clicking on a tag triggers the search mechanism for this tag, limiting its scope to annotations, thus retrieving the resources it has been used to annotate categorized by type and ranked by relevance (see Section 3). Additionally, an overall tag cloud is available, that aggregates terms and their frequencies across deliberations.

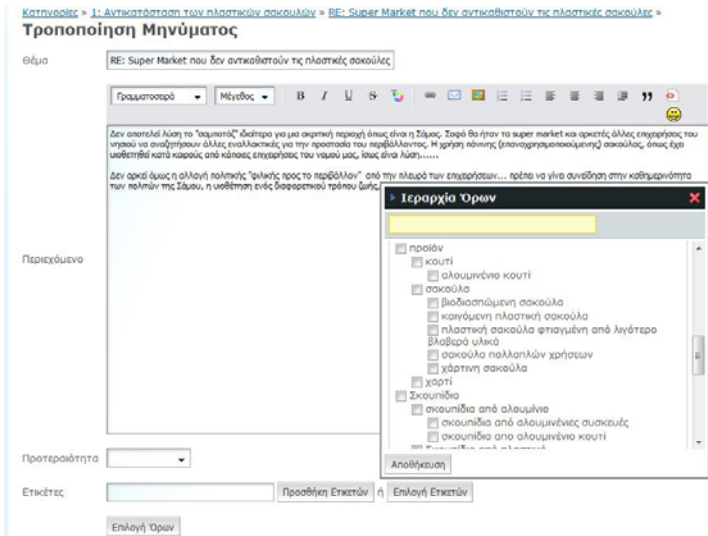


Fig. 2. Content annotation using tags

It must be emphasized however that despite the simplified view of the ontology that users are presented with, a connection of taxonomy terms, or tags, with their counterpart ontological terms is retained. Additionally to the use of a controlled vocabulary to annotate resources, custom tags may also be used in order to allow users to complement the vocabulary with their own terms.

3 Finding Relevant Views

Besides browsing via the tag cloud or navigating via the web front-end menu structure, we wanted to enable users to locate content quickly via an intuitive search mechanism, tailored to our eParticipation context and infrastructure requirements, namely:

- **Usability by general public.** The quintessence of the deliberation process is to involve as broad an audience as possible. Therefore, we had to take a 'least common denominator' approach, focusing on providing an intuitive search mechanism that performs well without presenting users with a complicated set of options.
- **Applicability to heterogeneous content.** In our platform, a variety of tools is used to support the deliberation process, each one generating content with different characteristics. Therefore, according to the 'least common denominator' approach we choose to rely on a combination of full-text search and user-provided annotation, as all the tools generate content that can be annotated, indexed and retrieved.

- **Applicability to user-generated content.** Finally, content is mostly generated by end users, therefore we cannot rely on resources to be properly annotated, even when using the simplest possible form of annotation (tagging). When present though, we want to take this annotation into account.

3.1 A Hybrid Search System

Taking the above points into account, we chose to pursue a hybrid search approach that relies mostly on full-text search and statistical measures to retrieve and rank results, but also uses annotations, where available, to influence ranking as well as domain ontologies to expand results.

Each piece of content generated by the platform tools is treated as a document and indexed using the Lucene engine². The title, tags and actual content of the piece are linguistically analyzed before they are indexed (as discrete fields), using a custom analyzer developed especially for the Greek language that performs tokenization and lemmatization. Additional fields are also used for some specific content types, based on their structure.

The search interface is extremely simple, having only one input field in which keyword search terms are entered (Figure 3). The query string is parsed using the same analyzer used for indexing, ensuring that query terms match index terms, and the resulting terms are looked up in the index. Ranking results is done using standard TF/IDF metrics, however the query performed against the index combines field scores giving different weights to each, promoting the tags field to reflect the fact that we consider explicit annotation as most relevant. Title field is weighted 2nd, content field 3rd and custom content fields follow.

This way we produce results that are ranked based on a mixed model, utilizing annotations, when present, while smoothly blending with an effective full-text search model. This takes place intuitively, by utilizing a different, simpler mechanism than the one utilized in similar existing work [4], as our context differs and hence so do our requirements. Additionally, queries can be further refined or expanded depending on user input.

Query refinement works on the deliberation level, so if users choose to use the search functionality while browsing a specific deliberation space, they can activate an option that restricts the results to content generated only for the deliberation at hand.

Query expansion works by utilizing the same search mechanism, but using a semantically expanded query as input. When the query expansion mechanism is activated, the original user query is first processed by the query expansion engine and then the resulting semantically expanded query is used as input for the search mechanism. Terms retrieved via the semantic expansion mechanism are displayed separately in a special 'expanded query tag cloud' that visualizes which are the extra terms and how frequently they occur. This is used as a means of filtering expanded results, as the user is not only presented with a list

² <http://lucene.apache.org/>




Fig. 3. Search Interface

of related terms, but is also able to preview how many results each of them is associated with. By clicking on a term the user is presented with related results.

3.2 Query Expansion via Contextualization

In order to perform semantic query expansion, we have implemented a component that works by determining for a given query term, terms that are semantically similar in the specific search context. The component works by utilizing domain ontologies and a corresponding semantic query expansion method [1]. Since this mechanism relies on the existence of a domain ontology, it only works for searching within deliberations that have been associated with one, via the process described in Section 2.

The method draws upon ideas and techniques from the areas of ontologies and fuzzy logic and allows for the customization of the semantic similarity assessment mechanism according to the search scenario’s particular characteristics. This is made possible through a context model that captures information about which of the domain ontology’s relations and to what extent should participate in the similarity assessment process. The parameters of this model are typically determined prior to the deployment of the system and are used by a contextualization algorithm for producing a term semantic similarity index (i.e. pairs of similar terms) that is subsequently used for the expansion [1].

The distinguishing characteristic of this expansion method compared to others is that it may utilize not only the *subclassOf*, *equivalent* and *instanceOf* relations that are typically found within an ontology but virtually any relation between concepts and their instances that has been defined within the domain ontology.

Thus, given a query, the platform’s query expansion mechanism works as follows:

- **Query parsing.** Query terms are parsed and compared against taxonomical terms generated via the association process described in Section 2. This is done using the same advanced analyzer used for indexing, thus giving great flexibility in the matching process, and results in a list of taxonomical terms.
- **Query expansion.** As noted in Section 2, taxonomical terms retain their association to their ontological counterparts, by storing their URIs. So at this point, the list of URIs that correspond to the taxonomical terms identified in the query parsing stage are sent to the query expansion component, which then uses the semantic similarity index described above in order to produce a list of ontological terms that are related to the ones given as input.
- **Query execution.** Finally, the list of ontological terms returned by the query expansion component are associated to their taxonomical counterparts (again, via their URIs) and used as input to the hybrid search mechanism. It must be noted that the weight of terms retrieved via the query expansion mechanism is appropriately reduced in order to influence their standing in the overall ranking and reflect the fact that they are deemed less relevant than terms contained in the original query.

4 Sharing Views

As eParticipation matures as a discipline, we expect to see more eParticipation projects being implemented and, consequently, the need to connect these communities and let them share their views to emerge. This is not just a hypothetical projection, as in addition to the City of Trikala, there are already similar projects under way in the City of Arta and the City of Thessaloniki, the 2nd largest in Greece. Although each community has its own characteristics and issues at stake, certain issues will be shared on the local, regional, national and international level, for which the exchange of experience shall be beneficial.

On the other hand, even though a specialized eParticipation framework best serves an integrated approach in a controlled environment, dialogue and argumentation exchange naturally also occur in non-designated environments and in non-prescribed forms. Fora, social networking sites, mailing lists and newsgroups alike, as different as they may be in terms of technical infrastructure, they all provide means to facilitate discussion for the user communities that use them. This is another 'resource', or content provider, that we would like to tap on by enabling exchange and reference to expressed views.

In order to address both of these challenges, we need to adopt an approach that enables us to not only specify the semantics of the exchange (what is the community and its discussions about) but also make it feasible technically, by providing a mechanism to support identification and cross-reference of resources, a transport layer and remote structured querying facilities. In that view, the adoption of Linked Data is a natural choice, as it covers all these requirements in a standardized way via ontologies, URIs and SPARQL [6].

In this section, we describe our approach towards the implementation of this vision, by describing our eDeliberation ontology, its use for making content

available as Linked Data as well as an application built on top of this infrastructure that showcases its potential.

4.1 The eDeliberation Ontology

The eDeliberation ontology is a conceptualization of our framework. Its primary goal is to make the structure and semantics of the framework explicit and subsequently to make the content it generates accessible to external applications in the form of Linked Data, rather than constitute a generic knowledge map of the eParticipation domain [14].

Each deliberation process is modelled within the eParticipation platform as an online community whose members make use of certain tools, such as fora, questionnaires and polls, in order to participate in each of the deliberation's stages. All the tools are associated with a community space in the platform. In this process, content such as forum posts, answers to questionnaires or documents is generated and annotated by means of the deliberation's associated domain ontology. This means that the deliberation ontology needs to represent each deliberation stage in connection to the associated tools and the content they generate. Additionally, it should provide a way to make content meaning explicit, by relating its tags to terms derived from the deliberation's domain ontology, as well as provide information about the use of tags in specific annotations.

In order to cover the above requirements, the eDeliberation ontology reuses well established ontologies from the Semantic Web community, namely the SIOC Core Ontology³, the Tag Ontology⁴ and the MOAT Ontology⁵, all suited for semantically representing Web 2.0 content. In particular, the SIOC ontology enables the integration of online community information and social web data, the Tag ontology models the relationship between resources, tags and users and the MOAT ontology enables users to define the meaning(s) of their tag(s) by using URIs of Semantic Web resources. Certain concepts and relations from these ontologies are either used as-is or extended in the eParticipation ontology.

Reused concepts are `sioc:Forum`, `sioc:Post`, `tag:RestrictedTagging`, `tag:Tag` and `moat:Meaning`. As for the other classes, these are:

- **Deliberation**: The process described in Figure 1. The concept is defined as a subclass of the concept `sioc:Space`.
- **TopicSelectionPoll**: Poll which is conducted for determining the topics that will be discussed in a deliberation (extends `sioc:t:Poll`).
- **PotentialDeliberationTopic**: Topic which has been suggested for deliberation.
- **Document**: A document that contains information regarding a deliberation.
- **Document Library**: The set of documents that regard a specific deliberation. The concept is defined as a subclass of the concept `Container` of the SIOC ontology.

The relations that link the above concepts are shown in Figure 4.

³ <http://sioc-project.org/>

⁴ <http://www.holygoat.co.uk/projects/tags/>

⁵ <http://moat-project.org/>

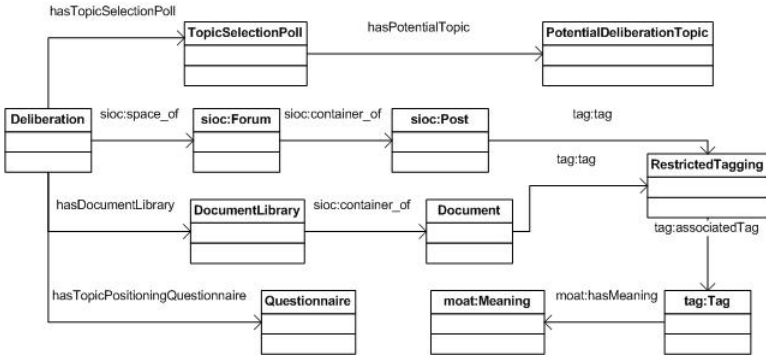


Fig. 4. Deliberation Ontology

4.2 Making Deliberation Content Available as Linked Data

Same as domain ontologies, the deliberation ontology serves a practical goal besides conceptually specifying the domain: it is used in order to make content generated on the platform available as Linked Data. We rely on existing work on the inbound/outbound Linked Data approach to make this possible [2].

As the eParticipation platform is implemented as a typical Web 2.0 n-tier application, all the entities and relations modelled in the eParticipation ontology live in a relational database. In order to make them available as Linked Data, we mapped the platform’s database to the ontology using D2R as the tool to produce the mappings as well as publish the data via a SPARQL endpoint. Most of the mapping process was rather straightforward, as the vocabularies used are a natural match for our domain the eParticipation ontology is also modeled appropriately.

One point that deserves special attention is the use of the Tag and MOAT vocabularies to represent both taxonomical and user-provided tags. In both cases, Tags and Taggings are mapped in the same way, however the difference is that taxonomical tags have exactly 1 pre-specified meaning, which is the same for every tagging and corresponds to the concept of the ontology from which they originate while user-provided tags may have from 0 to N meanings, depending on their use and to whether the user has grounded them with some ontology concept or not. In our case, user-provided tags do not have meanings, for reasons we shall explain in Section 6.

4.3 Distributed Contextual Views Retrieval

Having the ontological infrastructure and the Linked Data mechanism in place, we utilize them in order to provide context for discussions taking place in deliberations. The idea is to enable people participating in the discussion to have an idea of

what has been discussed elsewhere on the topic at hand. Each deliberation is associated with a forum that plays the role of a container of different discussion categories, which in turn may contain a number of discussion threads consisting of a series of messages. We want to locate discussions that have taken place on other sources and are deemed similar, based on some criteria, to the ones that are taking place in each deliberation and present them to forum users.

Design parameters. In order to do so, we need to define what it is we want to present and at what level of this discussion structure to associate it with and present it (*granularity*), where to retrieve from (*source selection*), how to retrieve it (*query mechanism*) and how to calculate similarity (*similarity measure*). We shall elaborate on the available choices for each of these parameters and how we define the positioning of our application on the selection space.

Regarding the choice of *query mechanism*, since we wanted to implement an application that does not rely on proprietary APIs in order to retrieve information from different sources, our choice was between using OpenSearch⁶ and Linked Data. We decided to go with Linked Data, as it allows for more flexibility in specifying data structure, semantics and search criteria. OpenSearch only supports text-based queries and is also not transparent as to how the results returned are ranked, while we wanted to have control over the ranking. Additionally, Linked Data is gaining traction, so we expect to see an increasing number of discussion board providers supporting it in the near future.

The *source selection* dimension however highlights a contradiction with regard to the use of Linked Data as the foundation for building applications. On the one hand, basing our retrieval mechanism on Linked Data means that, if properly designed, it can be generic enough to be applicable to any Linked Data source. In practice however, our choice is rather limited since at this point the use of Linked Data to publish discussion board data in structured form is still very limited. For this reason, we chose to only interlink different instances of our platform, each one serving as a data provider and a consumer at the same time.

In terms of *granularity*, the two extremes would be to establish connections at the deliberation level and the message level. Single messages taken out of context do not provide enough information to create associations, while deliberations on the other hand are eParticipation-specific constructs, so establishing connections at this level would mean that the mechanism would not be applicable to any discussion board. So we chose to focus our interlinking algorithm on retrieving threads of discussion related to a certain category, which is generic enough to work for any discussion board, while also utilizing eParticipation-specific information in case it is present.

Finally, given these choices, we designed a first version of the contextual views retrieval mechanism that calculates *similarity* using a mixture of structural information and content similarity on the textual level, making no assumptions about the presence of annotations, but is also generic enough so that it can be extended to take such information into account, if present.

⁶ <http://www.opensearch.org>

Implementation. Our distributed contextual views retrieval mechanism works as follows:

Source selection is currently controlled via a manually maintained list of SPARQL endpoints that provide discussion board data in a structured form, in compliance with the SIOC vocabulary. If the endpoints also provide deliberation data, then we assume that the endpoint is an instance of an eParticipation platform, so administrators are given the option to associate deliberations running locally with others residing on remote endpoints, using a special GUI to do so.

When presenting deliberation fora, users may choose to view relevant views for each discussion category by clicking the corresponding button. Then the name of the category is parsed (again using advanced language processing techniques) and the terms it contains are extracted and used to query endpoints, retrieve and rank relevant discussion threads and present links to them alongside the local threads.

There is a 'default' similarity measure used in case there is no annotation present, which is generic enough to be usable for any SIOC-enabled SPARQL endpoint: a SPARQL query is issued against the endpoint, searching for posts that contain at least one of the query terms in their content or title, or in their thread's title. This query returns information about Posts and their associated Threads that is processed locally to rank the results in terms of their syntactic relevance to the issued query terms, using information retrieval techniques similar to the ones used to rank local results.

Since this intermediate list of results contains discrete Posts and not Threads, we use it to produce a list of discrete Threads to be presented, based on: i) The sum of occurrences of Posts belonging to each Thread in the intermediate list. The more Posts each Thread has, the higher it is ranked. ii) Possible association between the deliberation to which the local discussion category belongs and the deliberation to which the remote Thread belongs. If such an association exists, the Thread is promoted in the ranking. Further elaboration is beyond the scope of this paper and is identified as future work.

5 Evaluation

At the time of writing, the system had been in operation for a period of one month, running a total of 4 deliberations. Some early indicators on which we can judge the platform's impact and importance are the following:

- Number of citizens who have registered at the platform: 70.
- Number of citizens who participated in the first agenda setting e-poll: 25.
- Number of citizens who posted to the e-forum: 20 with 33 posts.
- Number of e-surveys submitted: 4 answered in total by 44 citizens.
- Monthly average statistics: 784 visits, 385 unique visitors, 6.469 page views.

Although the above numbers may seem small, when evaluating them we have to consider the limited amount of time the project has been running, the novel

nature of the project and also the fact that in the Prefecture of Samos (with a population of 33.814 people), 42% of the population are internet users out of which 60% have broadband access⁷. Assuming that broadband users are the ones most likely to use such a sophisticated platform, then we can -very roughly- estimate that 1% of them became registered users and 10% participated in one way or another. We expect these figures to rise substantially in the coming months, as the recently launched publicity campaign for the project will reach its peak and word-of-mouth advertisement will also start taking effect.




Fig. 5. Qualitative user evaluation

Besides quantitative indicators provided by usage statistics, a user survey has also been undertaken in order to evaluate the system qualitatively. A total of 12 users so far have provided feedback by means of filling in questionnaires to evaluate the accuracy of the retrieved results as well as overall satisfaction by the system. The results of the evaluation are depicted in Figure 5, showing that the system scores well on ease of use and extremely well in the other categories.

6 Conclusions and Future Work

In this paper we described how Semantic Web technologies were applied in our eParticipation platform, deployed in the Prefecture of Samos. We gave some

⁷ Figures calculated and based on official data from the Observatory for the Information Society eEurope/i2010 July 08 report and the Census 2001 population data of the Greek National Statistics Office.

background on the domain and the platform and showed why the technology fits the domain, how it has been implemented using a Web 2.0 user-friendly approach as far as the user interface is concerned and what benefits it brings in terms of content annotation, browsing and search. We emphasized the sharing aspect of the platform, showed how it aligns with the Linked Data philosophy and infrastructure and described a distributed approach to contextual views retrieval that is based on it.

In terms of future work, we identify the following areas:

- **Improved support for annotation.** In order to assist users in the annotation process, we are looking to add a semi-automated annotation feature: the system will examine new content as well as the domain ontology in the context (deliberation) it has been generated and provide suggested taxonomical terms to use as annotations.
- **Incorporation of meanings for user-provided tags.** As part of an inbound/outbound Linked Data approach, we have been developing an application that grounds tag meanings with DBpedia concepts [2]. We shall include this application in the platform, as tag meanings can enhance search, disambiguate tag use and even be used for ontology alignment and evolution purposes, as soon as a Greek version of DBpedia is available.
- **Support for inference.** A known issue for the recipe we used to publish Linked Data is the lack of inference capabilities in the SPARQL endpoint provided by D2R [5], as 'a certain amount of inference (datatypes, basic RDFS, owl:sameAs) would be desirable and feasible'. Furthermore, by adding inference capabilities for each node to the eParticipation scenario we have, in which each community maintains control over its ontologies, it becomes an ideal match for a distributed reasoning scenario [3].
- **Refinement of similarity measure used in distributed contextual views assesment.** The similarity measure that we described in Section 4 is a first approach to assessing similarity between threaded discussions based on textual search. This has certain drawbacks, as on the one hand our D2R-based SPARQL endpoint does not currently support Semantic IR [13] features, therefore we have to rely on the limited matching capabilities offered by pure SPARQL, on the other hand we do not take into account the existence of annotation. While the former is beyond our control, for the latter we could utilize ontology alignment techniques and tools [8,7] in order to align ontologies used in different nodes and thus be able to assess similarity between resources annotated using these ontologies.

Concluding, we believe that the eParticipation domain is one that showcases very well the benefits of Semantic Web technology in action and we shall continue in pursuing this approach and further evolving it, in the hope that it can contribute to facilitating a dialogue ecosystem.

Acknowledgements

We would like to thank Mr. Vasilis Goulandris, for his contribution in the methodological part of the platform, the Prefecture of Samos Director, Mr. Manolis Karlas for his support in the implementation of the project, as well as the IT department of the Prefecture of Samos for their support during the platform's operation.

References

1. Alexopoulos, P., Wallace, M., Kafentzis, K.: A Fuzzy Ontology Framework for Customized Assessment of Semantic Similarity. In: 3rd International Workshop on Semantic Media and Adaptation (2008)
2. Anadiotis, G., Andriopoulos, P., Vekris, D., Zosakis, A.: Linked Data for the masses - using open source infrastructure and the inbound/outbound Linked Data approach to bring added value to end user applications. In: I-KNOW '09 and I-SEMANTICS '09 (to appear, 2009)
3. Anadiotis, G., Kotoulas, S., Siebes, R.: An Architecture for Peer-to-peer Reasoning. In: New Forms of Reasoning for the Semantic Web, CEUR Workshop Proceedings, vol. 291. CEUR-WS.org (2007)
4. Bhagdev, R., Chapman, S., Ciravegna, F., Lanfranchi, V., Petrelli, D.: Hybrid Search: Effectively Combining Keywords and Ontology-based Searches. In: Bechhofer, S., Hauswirth, M., Hoffmann, J., Koubarakis, M. (eds.) ESWC 2008. LNCS, vol. 5021, pp. 554–568. Springer, Heidelberg (2008)
5. Bizer, C., Cyganiak, R.: D2RQ - Lessons Learned. W3C Workshop on RDF Access to Relational Databases (October 2007)
6. Bizer, C., Heath, T., Idehen, K., Berners-Lee, T.: Linked data on the web. In: WWW '08: Proceedings of the 17th international conference on World Wide Web, pp. 1265–1266. ACM, New York (2008)
7. Bizer, C., Volz, J., Kobilarov, G., Gaedke, M.: Silk - A Link Discovery Framework for the Web of Data. In: 18th World Wide Web Conference (April 2009)
8. Euzenat, J., Shvaiko, P.: Ontology matching. Springer, Heidelberg (2007)
9. Kafentzis, K., Alexopoulos, P., Korakas, C., Goulandris, V., Baseas, I., Georgiolis, P.: Enabling citizen participation in local governance using ict. In: 3rd National Conference of the Hellenic Society for Systemic Studies, Piraeus, Greece (2007)
10. Kotis, K.: On Supporting HCOME-3O Ontology Argumentation using Semantic Wiki Technology. In: Meersman, R., Tari, Z., Herrero, P. (eds.) OTM-WS 2008. LNCS, vol. 5333, pp. 193–199. Springer, Heidelberg (2008)
11. Kotis, K., Vouros, G.: Human-Centered Ontology Engineering: the HCOME Methodology. *International Journal of Knowledge and Information Systems* 10, 109–131 (2006)
12. Macintosh, A.: eParticipation in Policy-making: the Research and the Challenges. In: Exploiting the Knowledge Economy: Issues, Applications, Case Studies, IOS Press, Amsterdam (2006)
13. Minack, E., Siberski, W., Nejd, W.: Benchmarking Fulltext Search Performance of RDF Stores. In: Aroyo, L., Traverso, P., Ciravegna, F., Cimiano, P., Heath, T., Hyvönen, E., Mizoguchi, R., Oren, E., Sabou, M., Simperl, E. (eds.) ESWC 2009. LNCS, vol. 5554, pp. 81–95. Springer, Heidelberg (2009)
14. Wimmer, M.: Ontology for an e-participation virtual resource centre. In: ICEGOV '07: Proceedings of the 1st international conference on Theory and practice of electronic governance, pp. 89–98. ACM, New York (2007)