Ontology-driven Collaborative Annotation in Shared Workspaces

Published version:

DOI:10.1016/j.future.2015.04.013

Ontology-driven Collaborative Annotation in Shared Workspaces

Anna Goy*, Diego Magro, Giovanna Petrone, Claudia Picardi, Marino Segnan

Dipartimento di Informatica, Università di Torino

C. Svizzera, 185, 10149 Torino, Italy

email: {annamaria.goy, diego.magro, giovanna.petrone, claudia.picardi, marino.segnan}@unito.it

Abstract

In the last decade, collaboration and sharing on the Web have become mainstream. Digital, remote interaction happens on a daily basis, not only to share digital resources, but also to create, manage and discuss them, in every possible situation were collaboration is required: from work teams to groups of friends, from community committees to no-profit organizations. In this paper we address the task of collaborative management of digital resources within a team, with a special focus on the task of semantic annotation, where team members, possibly supported by automated reasoning, enrich resources with properties that help in organizing, retrieving and creating connections between contents of different types. We focus in particular on the problem of reaching an agreement on the annotation itself among the participants. The paper presents a qualitative user study aimed at observing users behavior when faced with this task. The results of the study are then analyzed in order to draw guidelines, which are then implemented in a tool for collaborative annotation. This study is carried out in the context of the Semantic Table Plus Plus (Sem T++) Project, a framework supporting collaboration over thematic workspaces, whose goal is to enhance cooperation through awareness, enhanced communication and easy sharing of digital content.

Keywords: online collaboration, collaborative semantic annotation, ontologies, knowledge representation, knowledge sharing.

1. Introduction

The recent evolution of the web brought along new perspectives for its users. In particular, we embrace the idea that the current Web, sometimes referred to as *Web 3.0*, can be considered as the

* Corresponding author
Email address: annamaria.goy@unito.it
Postal address: Dipartimento di Informatica, Università di Torino, C. Svizzera, 185, 10149 Torino, Italy
Phone: +390116706786
Fax: +39011751603

evolution of "participative", service-oriented *Web 2.0* when it meets the *Semantic Web* (see [1] for an early and effective discussion about this point of view). Moreover, the convergence of social aspects and semantic technologies also led to a new vision which is referred to as *Social Semantic Web* [2]. In this perspective, the following aspects are important for the approach presented in this paper: both software applications and web sites tend to turn into interactive online services; emerging web technologies encourage participation and collaboration among users; the usage of semantically-based metadata about web resources is exploited in order to make them processable by "intelligent" software agents.

This scenario is experienced in our everyday life: in recent years, almost everybody has faced the growing need of interacting with online services and managing digital resources, in order to carry on many daily tasks, including work and business (e.g., managing projects, organizing and sharing digital documents, and so on), interaction with e-government applications, family management and leisure (e.g., online shopping or reservations for travels, hotels, restaurants, etc.). Moreover, users can take advantage from the enhanced possibility of sharing and collaboratively handling such digital resources, for example in social tagging systems, project management applications, collaborative editing tools, online conferences, file sharing tools, just to mention a few within a very large set.

In summary, the current web poses the challenge of managing a very large amount of heterogeneous, and often shared, digital resources (documents, pictures and videos, emails, posts, bookmarks, etc.), but it also suggests the techniques and tools to face such a challenge, by providing the set of technologies and approaches usually grouped under the label of *Semantic Web*. In particular, from the perspective sketched so far, the most important aspect is represented by *semantic annotation*, which enables software applications to access, manipulate, and in some sense "understand" the content of digital resources.

In this paper we present a framework supporting the collaborative management of shared digital resources, designed and developed within the *Semantic Table Plus Plus* (Sem T++) project. A major role, within this framework, is played by formal semantic representations of *information objects*, collaboratively built up by groups of users working together in collaborative thematic workspaces. In particular, in this paper, we will focus on this activity, discussing the results of a qualitative user study aimed at analyzing user requirements and defining the collaboration model supporting collaborative semantic annotation of shared digital resources.

The rest of the paper is organized as follows. Section 2 discusses the most relevant related work, taking into account different research fields, which provide the background of our approach, focusing, in particular, on the fields related to collaborative semantic annotation. Section 3 presents our previous work on Sem T++ by describing its main features. Section 4, which represents the original contribution of the paper, describes the role of semantic annotations of information objects in Sem T++. This section also contains the description of a qualitative user study about collaborative semantic annotation, a discussion of its results, and an explanation of how the model based on such a study has been implemented in Sem T++. Section 5 concludes the paper by summarizing its findings and open issues.

2. Related Work

2.1. Sem T++ General Background

A survey and a discussion of existing Web-based applications supporting collaboration, including groupware and project management tools or suites, can be found in [3], a previous paper introducing our approach.

Sem T++ is grounded into the critique to the so-called *desktop metaphor*, and in the approaches trying to replace it. A good overview of this issue can be found in [4], where an articulate discussion of the problem and a presentation of the proposals going beyond it can be found. The major critique to systems based on the desktop metaphor outlined by editors and contributors of the mentioned book is that such a metaphor, based on application-centered and file-centered models, fails in providing a support to the new needs emerging from the current ICT scenario, and in particular the need for an effective support to user collaboration, heterogeneous objects and multiple contexts management. Some of the alternatives to the desktop metaphor have influenced the design of Sem T++ more than others. In particular, Sem T++ proposes the metaphor of "table", as opposed to the "desktop" – where a person usually works alone – since tables are places where many people can sit in group, discuss, share resources, and work together. In this perspective the most relevant alternative approaches are Haystack [5] and the proposals grounded into Activity-Based Computing [6], [7]. Haystack [5] provides users with a flexible and personalized control over resource properties which help to organize them into coherent workspaces, referring to specific *tasks*, by providing support for uniform annotation, links to other resources, and retrieved. Activity-Based Computing [6], [7] is based on the assumption that the most important principle exploited to organize is not *application* or *file*, but *user activity*, which enables users to build a set of applications and documents related by a common context. A more detailed account of this literature is provided in [8].

An interesting work, which aims at integrating desktop-based user interfaces and technologies from the Semantic Web, is the *Semantic Desktop* initiative [9], mainly developed within the NEPOMUK project (nepomuk.semanticdesktop.org). The main goal of this project was to support collaboration among knowledge workers, through the integration of existing applications, achieved by the definition of an open-source framework for implementing semantic desktops, based on a set of ontologies. A proposal to connect the Semantic Desktop to the Web of Data, which inspired the future enhancement of Sem T++ with a connection to open datasets (mentioned in Section 4.1), is discussed in [10].

The exploitation of semantic knowledge to support users in collaborative resources management underlies many approaches. For example, in [11] the authors present an ontology to support media data management within the CineGrid Exchange network. [12] presents an extensible and domain-independent ontology-based architecture for data management systems, aimed at enabling the creation, storage, validation, query, and search of large amounts of data (and metadata) in heterogeneous formats.

In a slightly different perspective, with more emphasis on the formal description of interactions, instead of resources themselves, [13] presents a tool supporting the analysis of trends and patterns in collaboration activities taking place within multi-disciplinary design teams and relying on web-based heterogeneous collaborative applications. The described tool, TCN (Team Collaboration Networks), exploits an ontology providing the vocabulary for describing interactions between persons and/or information objects.

2.2. Sem T++ Annotation Background

Approaches and systems supporting resources annotation have been designed and developed within different research communities, with different purposes and characteristics. In this section, we will provide a non-exhaustive survey, focusing on the characteristics which enable us to define the notion of "annotation" suited to the approach presented in this paper.

In NLP-oriented annotation tools, "annotations" are typically labels associated to phrases within a document and usually rely on a predefined annotation schema. Such labels can be "semantic" (for example, when annotators identify Named Entities, such as geographic places, people, events, and so on); in this case they are usually linked to a semantic resource, such as an ontology (e.g., [14] and [15] among many others). In these tools, the annotated entity is typically a sub-part of a document (a word, a phrase, a sentence, a paragraph) and in their collaborative version they usually support collaboration within a relatively small group of annotators (see, for instance, the presentation of GATE Teamware – and a survey of related tools – in [16]; see also [17]). *Phrase Detectives* (anawiki.essex.ac.uk/phrasedetectives/) [18] represents an exception, where crowd collaboration is supported.

A similar notion of "semantic annotation" has been exploited in the Knowledge Management field (see, for instance, [19]), where annotations can link document parts (e.g., words or phrases) to instances in a semantic knowledge base – and thus indirectly to a set of classes in a domain ontology (for example, an annotation can link the word "Torino" to the instance *Turin*, and thus indirectly to the class *City*) – or directly to ontology classes (if an annotation directly links "Torino" to the class *City*). In this case, again, the collaboration takes place within a somehow closed community (typically, a company).

[19] also contains an interesting survey of annotation frameworks. A good survey of ontology-based annotation approaches, typically related to the Semantic Web vision, can be found in [20]. In many cases, ontologies provide the metadata structure, usually describing document features, such as *author*, *date*, *format*, etc. (e.g. Dublin Core: www.dublincore.org), or features characterizing organization or people web pages (e.g., FOAF: www.foaf-project.org). In other cases, annotation systems can exploit domain-dependent semantic resources, for example in the geographic domain (e.g., the Getty Thesaurus of Geographic Names: www.getty.edu/research/tools/vocabularies/tgn). In this last case, annotations tend to describe the document *content*, and semantic resources are used as controlled vocabularies, providing metadata values.

A huge impulse in this direction has been provided by the Open Linked Data paradigm (linkeddata.org), given the fact that most datasets refer to one or more ontologies, or "semantic" vocabularies (e.g., DBpedia: dbpedia.org, GeoNames: www.geonames.org); we do not enter into a detailed discussion of this trend, since it is outside the scope of the current paper.

A different set of tools providing support for annotation are those which enable users to add comments (e.g., sticky notes) to web pages or to digital documents in general. In this case, the annotated entity is the whole resource and the annotation usually consists in a free text, with no reference schema; see, for instance, Bounce (www.bounceapp.com), Diigo (www.diigo.com), My-Stickies (www.mystickies.com).

A huge field which has represented a significative reference for the design of the support to the annotation activity in Sem T++ is represented by social tagging systems. Such systems rely on a collaborative tagging activity, consisting in many users (typically an open community) adding tags (i.e., metadata) to shared resources (documents, pictures, videos, bookmarks, etc.); tags are typically free-text. Collaborative tagging can lead to the creation of folksonomies, i.e., bottom-up classification models incrementally built by users themselves [21], [2]. Tagging systems and folksonomies have been largely considered as a way to overcome the rigidity of taxonomic classification, typically used in folder hierarchies, since they provide a multi-facets classification [22]. However, some works have shown how collaborative tagging systems share some drawbacks with taxonomies [23], and suffer from problems which are typical of natural language interpretation, such as polysemy, synonymy, ambiguity, and the "basic level" problem [22] [24]. One of the most interesting aspect of collaborative tagging systems is the fact that tags are used to describe very different aspects of the tagged resource, linked to user goals; for example, tags can be used to describe the resource *content* (when a document about New York is tagged with "New York"), to identify the resource *type* (when a web page is tagged as "article"), to state the author or owner (when a blog post is tagged with "Anna"), to provide an *opinion* on the resource (when a document is tagged as "useless", or "funny"), or to add information about a *task* referring to the resource (when an article is tagged as "to read") [24]. We will see in Section 4.1 how Sem T++ takes these aspects into account.

Recently, many approaches have tried to enhance tagging systems with semantic capabilities (e.g., [25]). For example, in [26] tags refer to a formal semantic knowledge representation, constraining – and at the same time empowering – the annotation. Also semantically-enhanced tagging systems are crowd-based, since annotators are simply web users, and usually do not belong to a small group or a closed community.

Annotations are also widely used in e-learning (e.g., see [27]): in this case, annotations can be more structured (e.g., they can refer to annotation models such as Dublin Core: www.dublincore.org) and can be attached to sub-parts of documents (e.g., chapters of a book). In e-learning, collaboration takes place typically within a community, possibly large, but which does not coincide with the indistinct crowd.

Another field where semantic annotations have been largely exploited are semantic wikis. Buffa and colleagues distinguish two approaches: "the use of wikis for ontologies" and "the use of ontologies for wikis" [28, p.2]. From our point of view, the interesting approach is the latter, which is also the perspective of the system proposed by the authors, SweetWiki, a wiki tool enabling a structured, semantic annotation of resources. An interesting aspect of SweetWiki is "the use of a 'wiki object model', an ontology of the wiki itself", which enables queries such as "show pages and videos that talk about this subject" [28, p.2]. SweetWiki object model shares with the semantic model of Sem T++ the idea of modeling aspects other than those representing what the domain resources "talk about" (see Section 4.1); however, differently from our system, it is oriented to social tagging, representing another example of semantic-enabled social system.

Since the focus of this paper is user collaboration in handling semantic descriptions of shared digital resources – which can be considered a collaborative semantic annotation activity – we conclude this section with a brief discussion of collaboration models used in annotation tools. NLP-oriented annotation tools usually include a support for user *roles* (e.g., managers, editors, annotators), and guide the annotation process through workflows (e.g., GATE Teamware [16]). Workflows, defining sequences of tasks related to the annotation process and assigned to annotators, are used also in AlvisAE [29]. Also collaborative annotation in e-learning systems, such as PAMS 2.0 [27], is based on the definition of a predefined set of user *roles*, coupled with group membership for controlling the access to resources and annotations. In SYNC3 [17] "no conflict resolution is performed: if two users annotate the same text segment, both annotations are kept into the system, no matter if they are overlapping or contradicting"; moreover, the Ellogon platform [30] provides "comparison facilities, to identify mismatches among independent annotations of the same document, or calculate inter-annotation agreement".

3. Personal Information Management in the *Semantic Table Plus Plus* Project

3.1. Overview

In Section 1 we claimed that users need to manage an everyday increasing amount of digital resources. Moreover, such resources usually belong to heterogeneous types (documents, emails, bookmarks, multimedia items, and so on) and, as a consequence, they are often encoded in different formats, handled by different applications, and stored in different places. This poses a great challenge for the individual user, who has to switch between applications and storage systems even when handling resources which refer to the same thematic context or activity. For instance, it is a very common experience dealing with a folder with documents, a list of bookmarks, many emails and possibly some videos, all concerning the same research project, or a planned journey. Typically, to access all these resources, we have to run different applications or services (e.g., a local or online word processor, a bookmark management application, an email client or a web mail service, and so on).

The Sem T++ project was born to face this basic problem: the idea is that the user should switch between different thematic contexts, possibly linked to specific activities, but she should be spared "useless switches", i.e. moving from the online word processor folder, to the bookmark managing system, or to the email client, in order to find a resource talking about a given topic, e.g., the project budget discussed in the last meeting or the opening hours of the MoMA Museum in New York; [3] and [8] describe the first version of the framework (i.e., T++), which did not include the semantic model supporting annotations.

Sem T++, like its previous non-semantic version T++, is based on the metaphor of *tables*, i.e., workspaces representing thematic contexts, usually linked to specific user activities. Tables are user-defined at the granularity level chosen by their creator (e.g., a table can be used to manage a work project, to plan a journey, to handle children care). Tables are populated with *objects*, i.e., abstract views over digital resources of different type, lying on the table and accessible from the point of view of their *content*. A table concerning the planning of a journey to New York will contain documents, bookmarks, emails, images, etc. all accessible from the table itself, all concerning travelling to New York, and each one talking about some specific information (the MoMA Museum, flights to NY, restaurants reviews, and so on).

Table objects are *annotated* in a uniform way: both *comments* and *semantic annotations* are available in Sem T++, and are homogeneously handled as referring to a given object, independently from the object type.

Tables are *shared spaces*: although, in principle, a user could create a "solipsistic" table, tables in Sem T++ have been conceived as intrinsically collaborative environments. The most basic support to users collaborating on a table is provided by workspace awareness tools, i.e.: (i) Standard awareness techniques, such as icon highlighting, are used to notify users about table events (e.g., an object has been modified). (ii) Tables implement *selective presence*: on each table, a presence panel shows the list of table participants, highlighting who is currently sitting at the table; when a user is sitting at a table, she is (by default) "inactive" at other tables. (iii) Notification messages, coming from other tables or even from outside the Sem T++ environment, are filtered on the basis of the topic context represented by the active table (see [31] for a presentation of notification filtering techniques).

As far as editable objects are concerned, we rely on existing applications (e.g., Google Docs), which typically handle issues related to user collaboration with specific policies [32]. Finally, semantic annotation of table objects, which represents the main topic of this paper, is a collaborative activity, and it will be discussed in detail in Section 4.

Fig. 1 shows the user interface, as implemented in the proof-of-concept prototype described in [8]. In the top-right corner, a panel displays the list of tables (only one in the displayed screenshot). On the left, users can find the main tools available on the current table: the access to a user interface for objects selection, the two basic communication services – the table *Chat* and the *Blackboard* (for sharing synchronous and asynchronous messages among table participants) – and a shared *Calendar*. In the bottom-right corner participants at the table in focus are listed. The central area is available for containing a panel representing the user interface of the functionality in focus (e.g., objects selection, shared calendar, and so on).

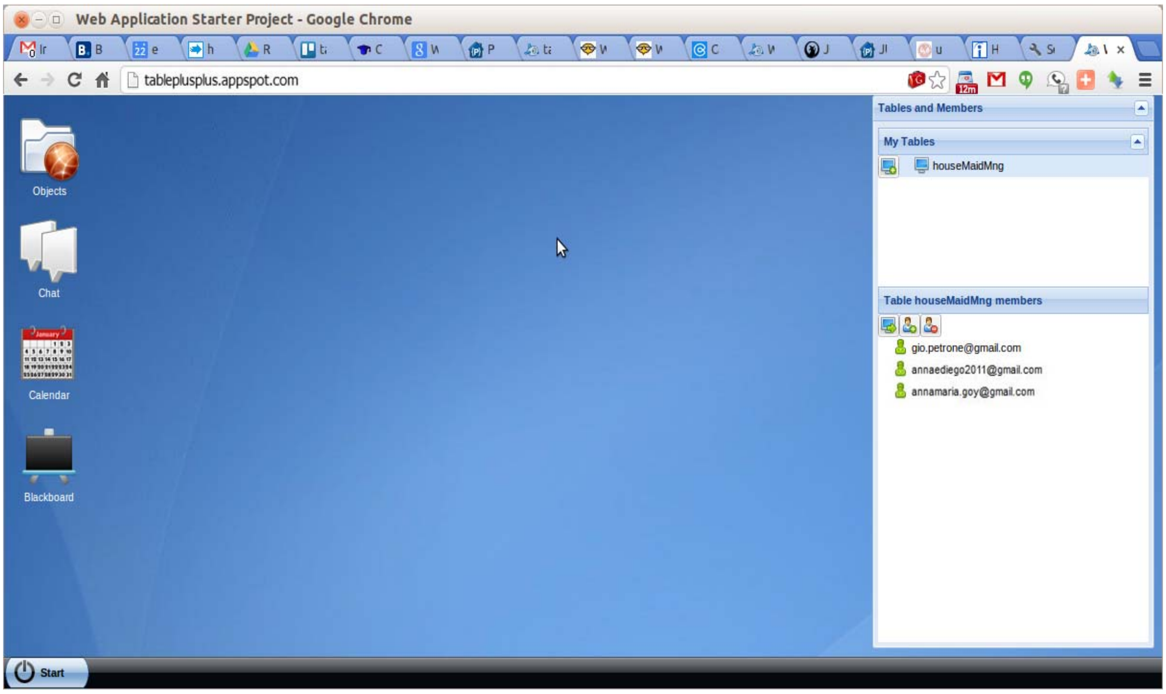

Fig. 1: Sem T++ proof-of-concept prototype user interface

In order to make tables "smarter", we provided them with knowledge about the objects they "contain", in the form of a *Table Ontology* modeling *information objects* with their properties and relations.

The Table Ontology provides the means to model, for example, the fact that a table object has parts (e.g., a document containing images and hyperlinks), the fact that it is written in French, or in Portuguese (or that it has parts written in French and parts written in Portuguese), the fact that it is encoded in PDF, the fact that it has an author, a main topic and its content refers to a set of entities (e.g., a web site presenting New York – its topic – and describing tourist attractions like the MoMA, Central Park, and so on).

The Table Ontology is an extension of the Knowledge Module of O-CREAM-v2 [36], a core reference ontology for the Customer Relationship Management domain, based on the foundational ontology DOLCE (Descriptive Ontology for Linguistic and Cognitive Engineering: www.ontologydesignpatterns.org/ont/dul/DUL.owl) [34], and one of its extensions – i.e., OIO (Ontology of Information Objects: www.ontologydesignpatterns.org/ont/dul/IOLite.owl) [35]. For a formal and detailed account of Sem T++ Table Ontology, and the semantic representations it supports, refer to [33]; here we briefly summarize the most relevant aspects.

Objects lying on tables are represented as instances of one of the subclasses of *InformationElement* (the "root" class of the Table Ontology) – i.e., *Document*, *Image*, *Video*, *EmailThread*, etc. – which inherits from the parent class several properties, among which, those representing:

- the relation between a resource (x) and its main topic (y), at time t[1] – *hasTopic*(x, y, t);
- the relation between a resource (x) and the entities it refers to, or it "talks about" (y), at time t – *hasObjectOfDiscourse*(x, y, t);
- the relation between a resource (x) and its authors (y), at time t – *hasAuthor*(x, y, t);
- the relation between a resource (x) and the languages used in it (y), at time t – *specifiedIn*(x, y, t);
- the relation between a resource (x) and the information objects (y) it contains (e.g., hyperlinks or images), at time t – *DOLCE* : *part*(x, y, t).

Moreover, in the Table Ontology, there are other properties used to describe the relations between a resource and the various formats used for its encoding, a property to link web resources to their URL, and some properties used to represent the relation between hyperlinks and resources they point to. Some examples of such properties are depicted in Fig. 2, where the (simplified) semantic representation of a table object (actually, a web page, encoded in HTML5, talking about New York and its tourist attractions, written in English by NY City Municipality, and containing several hyperlinks) is shown.

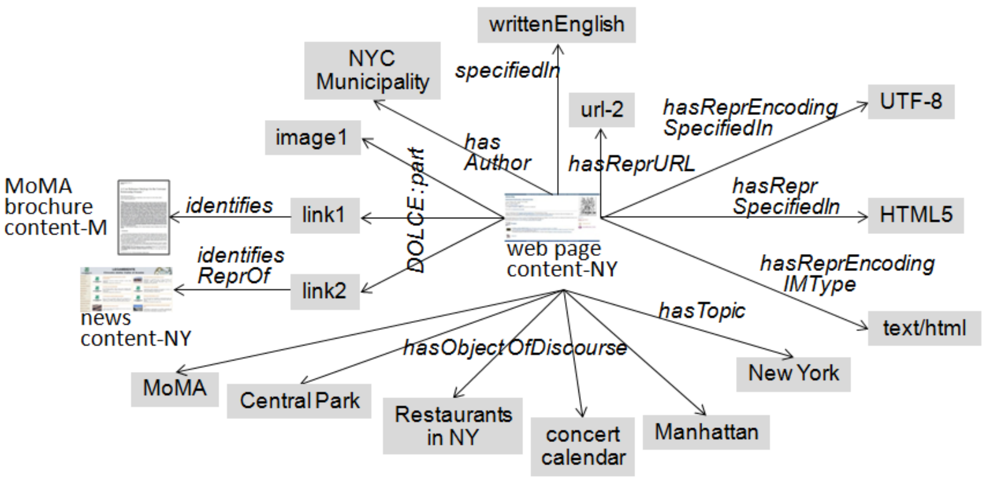

Fig. 2: The (simplified) semantic representation of a table object

The main advantage that "smart" tables, i.e., tables having an internal formal semantic representation of the resources they host, offer to Sem T++ users is a flexible access to digital resources, based on different criteria, which can be selected and combined by users themselves on the basis of their needs. For example, imagine that a user sits at a table about the activities of a NGO for environment safeguard and needs resources in order to write an article about waste

[1] Relations in the First Order Logic formalization of the Table Ontology contain a time parameter t. The current version of the system does not take the time dimension into account: this enables us to adopt a common approach in OWL ontologies, i.e., omitting the time parameter t, thus simplifying ternary predicates into binary ones.

recycling. She could start by selecting *waste recycling* from the list of topics present on the table, thus getting all table objects having that topic. Then she could get email conversations on the same topic, by filtering the previous list specifying *email* as object type, or she could get "political" opinions about the issue by filtering the object list specifying local political leaders as authors. The mechanisms supporting such an interaction are described in details in [33].

3.2. Architecture, Proof-of-concept Prototype and Evaluations

Fig. 3 shows Sem T++ architecture.

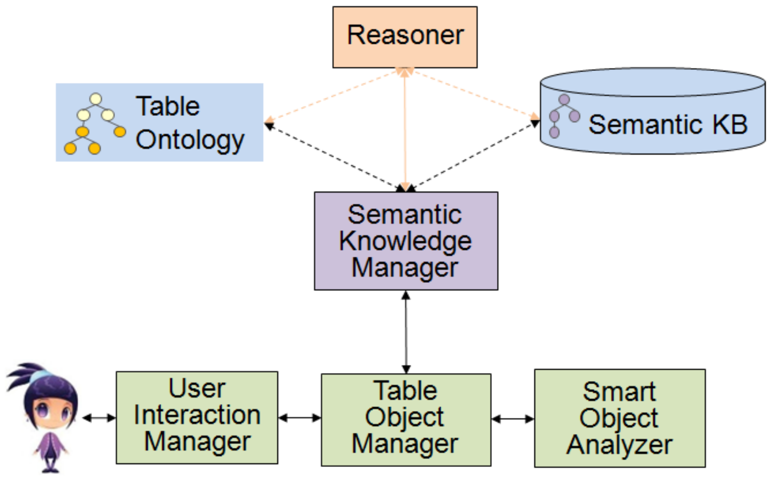

Fig. 3: Sem T++ architecture

The **Table Object Manager** is in charge of all the operations which take place on tables (e.g., adding/deleting resources, adding comments to a table object, and so on), while the **User Interaction Manager** handles all tasks related to the interaction with users. The **Smart Object Analyzer** provides the Table Object Manager with the analysis of table objects; for example, it detects parts included in the analyzed resource (e.g., images and hyperlinks), the language used, the encoding formats, and so on.

The **Semantic Knowledge Manager** manages the semantic model described above: its (static) knowledge about information objects is represented by the Table Ontology (see Section 3.1); moreover, it creates and modifies the semantic descriptions of table objects, which are stored in the Semantic KB, and invokes the Reasoner, when required (the role of the Reasoner in building semantic descriptions of table objects will be discussed in Section 4.1).

Sem T++ has been implemented in a proof-of-concept prototype in which the backend is a cloud application (a Java Web App deployed on the Google App Engine) accessible through a web browser; it relies on the API provided by Dropbox and Google Drive to store files corresponding to table objects and on Google Mail to manage email conversations. The Semantic Knowledge Manager (SKMgr), the Smart Object Analyzer (SOA), and the Table Object Manager (TOMgr) are all Java components. The Table Ontology and the Semantic KB are written in OWL (www.w3.org/TR/owl-features), while the Reasoner is based on Fact++ (owl.man.ac.uk/factplusplus). The SKMgr exploits the OWL API library (owlapi.sourceforge.net) to access semantic knowledge (i.e., the Table Ontology and the Semantic KB), while the SOA relies on a parsing Web Service, written in Python, to analyze resources. The User Interaction Manager has been implemented by a dynamic, responsive single page, exploiting AJAX to connect to a set of Java Servlets and exchanging JSON objects; its responsiveness is supported by using Bootstrap (getbootstrap.com). Most of the components of the prototype are heterogeneous, written in different languages and running on various operating systems, therefore they are implemented as RESTful Web Services exchanging JSON objects.

Some important functionalities of Sem T++ have been evaluated with users. In particular: [8] reports the results of an experiment in which participants were asked to perform a pre-defined sequence of collaborative tasks (including communication, resource sharing, and shared resources retrieval) using standard tools (such as email, Dropbox, Google Drive, Skype) and using T++ (the first version of Sem T++, without semantics). The results showed that user satisfaction is higher and performing the required tasks is faster using the T++ environment. [33] discusses the results of another user evaluation, in which potential Sem T++ users answered a post-test questionnaire after a guided interaction with Sem T++ UI mockup, in which they had to perform complex resource selection on a table. The results of the questionnaire confirmed that the availability of a user interface enabling the uniform and consistent management of different types of objects and the possibility of selecting table objects by combining different criteria (among which resources content) are highly appreciated features, which contribute to provide a better, less fragmented user experience and a more effective access to shared digital resources.

4. Sem T++ in a Collaborative Perspective: Semantic Annotation of Shared Digital Resources

4.1. Semantic Annotations

In Section 3 we mentioned the fact that Sem T++ has been conceived as a collaborative environment, and therefore its tables are fundamentally shared spaces. We also described why tables can be considered "smart", since they are endowed with semantic capabilities, based on formal descriptions of table objects properties. But how are such semantic descriptions built and modified by table participants?

In order to answer, we have to start from another question: **What are semantic annotations in Sem T++?**

Sem T++ shares features with several models underlying annotation systems, but it also has some peculiar characteristics. In particular, in Sem T++ the annotated entity is always the whole resource (i.e., a table object), and the Table Ontology acts as an annotation schema, by providing the structure of the annotation, i.e., the set of properties (*hasTopic*, *hasAuthor*, *specifiedIn*, etc., as described above) for which table participants can specify values. These properties represent a formal and well-grounded model reflecting the typical usage of tags to describe different aspects (content, resource type, author, ...) of the tagged resource, as discussed in Section 2.

The Sem T++ semantic model provides a semantically rich controlled vocabulary for the values of some properties (e.g., a list of natural languages, encoded as instances of the *NaturalLanguage* class in the Table Ontology for the property *specifiedIn*). However, for the properties representing the *content* of table resources – i.e., *hasTopic* and *hasObjectOfDiscourse* – the current version of the system does not provide any link to domain-dependent semantic resources (i.e., if a web page is about New York, no link to any instance representing *New York*, or to any class representing the concept of *City* in some ontology is currently provided). We are working at an enhancement of Sem T++ in this direction, by providing connections between topics/objects of discourse and external semantic resources, which can be found in open datasets such as DBpedia (dbpedia.org) or GeoNames (www.geonames.org).

From the point of view of user collaboration, in Sem T++, it takes place among table participants, i.e., pre-defined, typically small, groups of people, who have been invited to join a table with the goal of collaborating in a specific activity.

With all these characteristics in mind, we can now come back to the initial question: **How are semantic descriptions built and modified by table participants?**

When a new object is added to a table (or when an existing editable one is modified, for instance because a table participant adds an image in it), Sem T++ builds (or updates) the corresponding semantic representation (see the example in Fig. 2). Several property values are automatically set by the system, while for some others a decision by table participants is needed. In particular:

- Property values automatically determined by the system, relying on the analysis of the Smart Object Analyzer, are: partonomic relations, together with part types(classified as instances of, e.g., *Image, WebResourceLink*, etc.), and all format-related properties.
- Property values proposed by the system as *candidates* to table participants, who are asked to confirm or reject them and to possibly add new values, are found on the basis of:
    - The Smart Object Analyzer analysis of the resource, as far as authors and languages are concerned;

- The Semantic Knowledge Manager invoking the Reasoner, as regards – again – languages and authors[2], together with topics and objects of discourse.

Fig. 4 shows the prototype user interface enabling table participants to select objects of discourse from the set of candidates proposed by the system. The user can select candidate values by clicking on them in the *Candidates* panel.

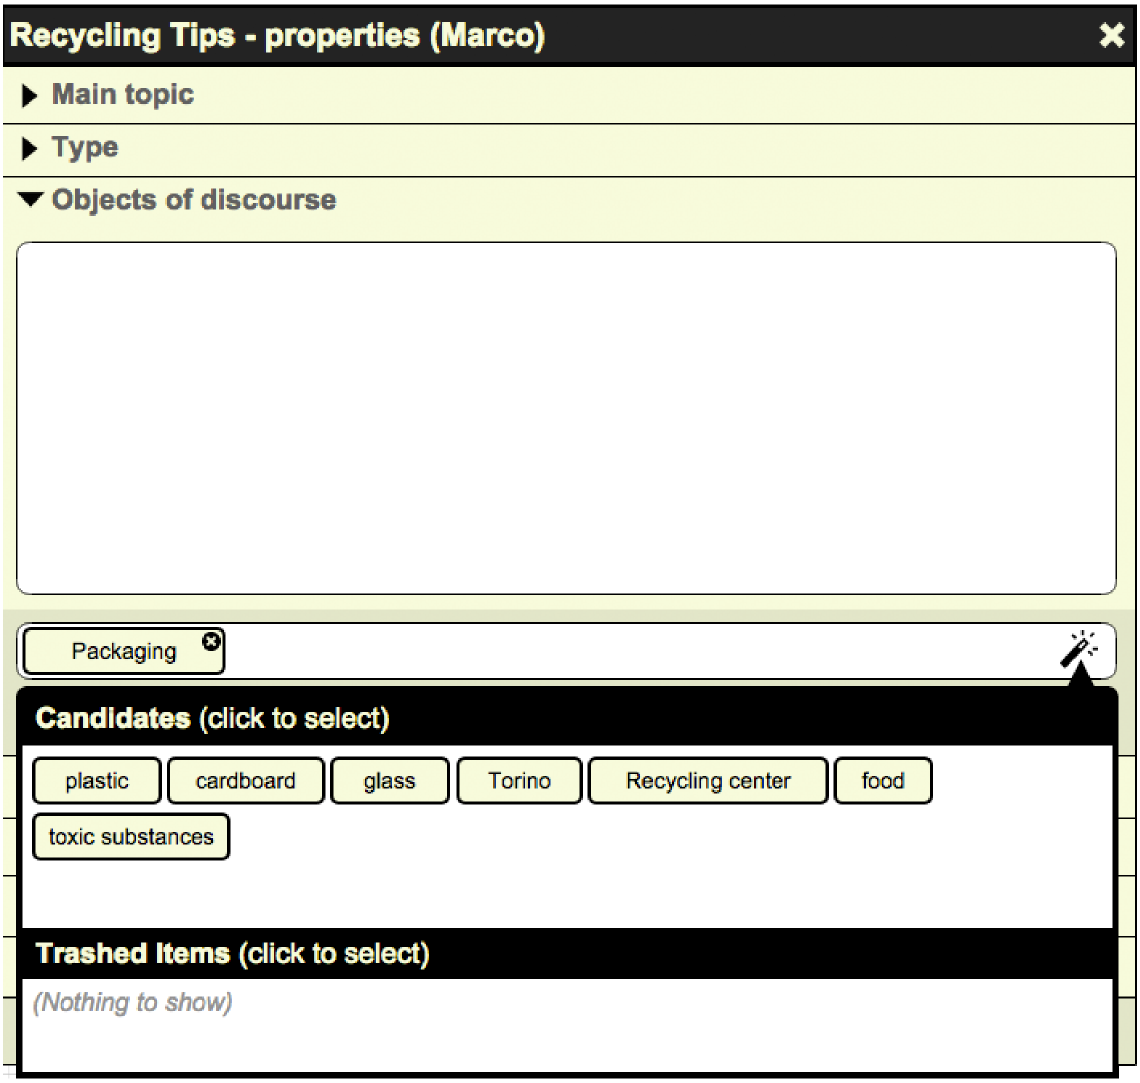

Fig. 4: Sem T++ user interface: the panel for selecting *candidate* objects of discourse

The Smart Object Analyzer detects the author and the language of a document by analyzing meta-data information when available (e.g., the *author* and *language* meta-tag in the *head* of an HTML document). The Reasoner, invoked by the Semantic Knowledge Manager, infers candidate values for the previously mentioned properties (language, author, topic and objects of discourse), thanks to a set of axioms, included in the Table Ontology on this purpose. We show an example of such axioms – i.e., the one concerning the *hasTopic* relation:

[2]The pool of candidates, concerning authors and languages, proposed to users, is the merge of the sets of candidates proposed according to Smart Object Analyzer analysis and to the inferences drawn by the Reasoner.

$$InformationElement(x) \wedge DOLCE : part(x, z, t) \wedge hasTopic(z, y, t) \rightarrow hasCandidateTopic(x, y, t)$$

For instance, the Reasoner can infer that *Central Park* (*y*) is a candidate topic of a document (*x*) from the fact that the document itself (*x*) includes an image (*z*) which has *Central Park* (*y*) as topic. If the user confirms a candidate value, the Semantic Knowledge Manager adds a new fact to the Semantic KB, stating the relation (e.g., *hasTopic(document, Central_Park, t)*).Obviously, the user can also reject a system suggestion and specify, for any property, a different value.

Now, **what happens in the process of building semantic representations of table objects if we shift the focus from an individual perspective to a collaborative one?**
Some of the properties defined in the Table Ontology should not pose any challenge, under this point of view: some of them are simply detected and set by the system (e.g., formats), while for others (e.g., the natural languages used in document or in a video) it is quite difficult to imagine table participants discussing or disagreeing about them. The properties which pose an interesting challenge with respect to user collaboration and knowledge sharing are those related to the *content* of table resources, i.e., the properties representing "what a resource is about": *hasTopic* and *hasObjectOfDiscourse*. The values assigned to these properties, in fact, depend significantly on the personal view each user has on the content (and scope) of table resources. Reaching a consensus on the annotation may therefore not be straightforward. Moreover, tables can be used to manage activities of any kind, including ones which probably require someone "having the last word" (e.g., a project supervisor). Thus, how do table participants collaborate on this activity? Is it easy or difficult to understand other people's contributions, and how can additional tools improve this understanding? What can help participants to converge toward a shared perspective? How does the chosen policy impact on the collaboration and on reaching a satisfactory result? In order to gain more insight on these issues, we designed a qualitative user study, described in the following sections.

4.2. Qualitative User Study

Our goal was to virtually observe some instances of the collaborative annotation process.
In particular, we wanted to observe the attitude of users when annotating resources:

- on tables created to manage different activities related to different contexts (e.g. workplace, family, friends, etc.);
- of different types (e.g. text, image, video, etc.);
- with different collaboration/decision policies.

As we already discussed, collaborative annotation presents some analogies with social tagging; to the best of our knowledge, however, no social tagging system allows to choose the collaboration policy. For this reason we were particularly interested in seeing if there were policies that were considered unacceptable, useless or annoying by our users, as well as whether some collaboration policy could lead to conflicts or – in the case of an "anarchic" policy – to the impossibility of reaching a stable version of the annotations.

4.2.1. Methodology

We recruited the users (15 people) among our friends on Facebook, according to availability sampling. Since our goal was not an extensive experimentation, but rather a thorough observation of small groups representing potential table participants in Sem T++, we did not aim at a representative statistical sample. By "potential table participant" we mean someone who may consider – and even welcome – the possibility of long-distance collaboration, not only in writing documents, but also in the collection, sharing and management of digital resources that are pertinent to the collaboration theme. Thus, our call was for people who had some experience with online group communication and with the type of cooperation and sharing that characterizes Web 2.0 applications. Also, we selected only people that could connect at least twice a day to the shared documents, in order to carry out some simple tasks.

We built 3 groups of 5 people (in the following labeled as **G1**, **G2** and **G3**). Each group was asked to participate in two scenarios where they were asked to collaborate in annotating two resources, chosen among several ones, providing for each of the two a *topic* and several *objects of discourse*. The presentation of each scenario and the list of resources were introduced in a Google Docs document. Each resource was represented by a hyperlink that users were asked to follow in order to read/see the content; all resources were non-editable. Users could provide their annotations in the same document, in a highlighted area below the hyperlinks pointing to the two resources they had to annotate.

The scenario presentation contained the following information:

- A description of the context for which the table was created. We imagined two different contexts: a collaboration on a project for promoting waste recycling (henceforth **REC**), and a collaboration concerning the creation and staging of an amateur play with disabled people, drawn from a children short story (henceforth **PLAY**).
- A list of the group members; as mentioned above we created three groups **G1**, **G2** and **G3**, and each group played both situations, so we had six scenarios.

- A description of the activity that they were supposed to perform, namely, annotation. Since the people we recruited did not know about our research work, and in some case were not familiar with online collaboration, we explained that they were supposed to collaboratively decide one "main topic" label (our *hasTopic* property) and several "characterizing" labels (our *hasObjectOfDiscourse* property) with the goal of simplifying future selection and retrieval of the resource on the table.
- Some hint on how to approach the task, namely that they could communicate with the other group members, to facilitate the process, by exploiting different means (both internal and external to Google Docs), and that they could feel free to delete or edit what other people did, since Google Docs keeps a history of all edits.
- A description of the collaboration policy; each group played with two of the following policies:[3]
    - **Consensual**, where the annotation process is closed only when consensus is reached among all participants. This policy was played by group **G1** in the **PLAY** situation and by group **G2** in the **REC** situation.
    - **Supervised**, where the annotation process is closed by the table supervisor (in this case we played this role ourselves) who may or may not participate in the annotation process, but has the final say on the chosen annotations. This policy was played by group **G1** in the **REC** situation and by group **G3** in the **PLAY** situation.
    - **Authored**, where the annotation process is closed by the resource creator, who has the last word on the chosen annotations. In this case a member of the group was introduced as the resource creator, and then privately contacted to receive specific instructions. This policy was played by group **G2** in the **PLAY** situation and by group **G3** in the **REC** situation.[4]
- A list of hyperlinks pointing to the shared resources, the first two of which had to be annotated. In the **REC** situation we asked to annotate an infographics (depicting common

[3] We did not ask our subjects to play out all three policies as we felt it was too burdensome and most people would have backed out. It was not easy to find volunteers who could be more or less connected for two consecutive weeks and have time each day to contribute. On the other hand, we wanted each group to play at least two policies, for several reasons: (i) to make them experientially aware of the policy issue; (ii) to be able to recognize group patterns that were independent from the policy; (iii) to be able to see the effects of a same policy in at least two different situations. This last point could have been addressed having more groups, but this would have been problematic due to our difficulty in recruiting.

[4] The Authored policy may seem similar to the Supervised one; however there are two key differences that can affect the collaboration process. First, the supervisor may not work alongside the team in the annotation phase while the author is actively participating (since he or she produced the resource). Another difference is that every team member can be the author of one or the other resource; therefore, even if on a single resource the policy looks similar to the Supervised one, if we take into account the collaboration process as a whole, we find no recognized leader but rather a sharing of responsibilities.

recycling mistakes) and a text (describing "what goes where when separating waste"), while in the **PLAY** situation we asked to annotate a text (the synopsis of the short story) and a video (an excerpt from a cartoon drawn from the same short story).

After a group had been assigned a situation to play, they were free to work for at least 2 days. We sent them every day an email reminder with a link to the Google Docs document to work on, asking to connect at least twice a day to check what others had done and possibly edit the annotations.
We monitored the situation by checking the revision history, to see whether people kept on changing annotations or we could safely presume they were satisfied with their work. At that point we closed the annotation phase by making the annotation document non-editable by group members and:

- in the **Consensual** policy, "freezing" the last version of the annotations;
- in the **Supervised** policy, selecting ourselves a final set of annotations – possibly undoing some of the last changes and/or inserting some brand new annotations;
- in the **Authored** policy, asking the fictitious resource creator to make a final choice of annotations according to her personal opinion.

After closing the annotation phase we asked the group to check the final annotations and see whether they were satisfied with them.
When a group had completed both situations, members were invited to fill in a questionnaire where they could answer some questions and provide observations on the collaborative annotation experience.

4.2.2. Survey Results

The questionnaire we presented to our users was organized into four sections. The first one aimed at profiling the users. We asked the subjects to self-assess their familiarity with web tools on a 5-point scale. Most people declared a good or very good familiarity (**5** people answered "4"and **9** people answered "5"; only one person answered "3"and no one answered "2"or "1"). **12** people had already used Google Docs, and **8** of them had used other collaborative tools. **3** people had never used Google Docs nor any other collaborative tool.
The second section of the survey inquired about the overall experience with collaborative annotation. People participated quite actively: on the average, they actively modified the annotations 2.4 times for each scenario. As shown in Fig. 5, most people found the experience quite **easy**, and reasonably **useful** (only 1 person gave an answer lower than 3). Opinions were more

divided on whether it was **funny** (where we see a perfect balance between "Boring" and "Fun") or **interesting** (where the scale is slightly tipped in favor of "Interesting").

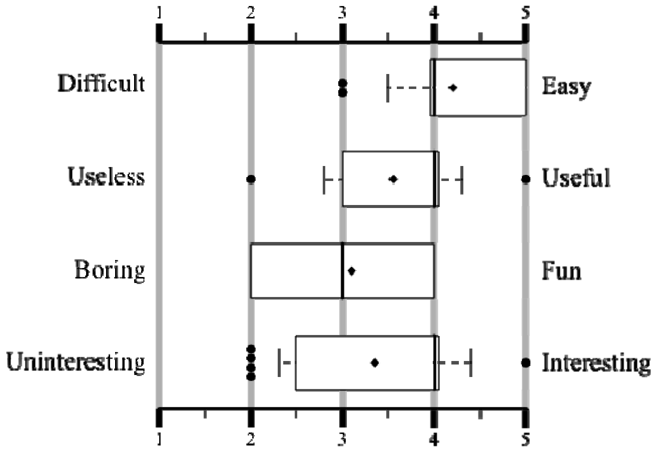

Fig. 5: Assessment of the quality of the collaborative activity. Each boxplot shows, besides the I, II and III quartiles, the standard deviation (whiskers), outliers (black dots) and the mean (diamond)

We also inquired how much they were satisfied with the final annotations. The results are shown, by means of boxplots, in Fig. 6.

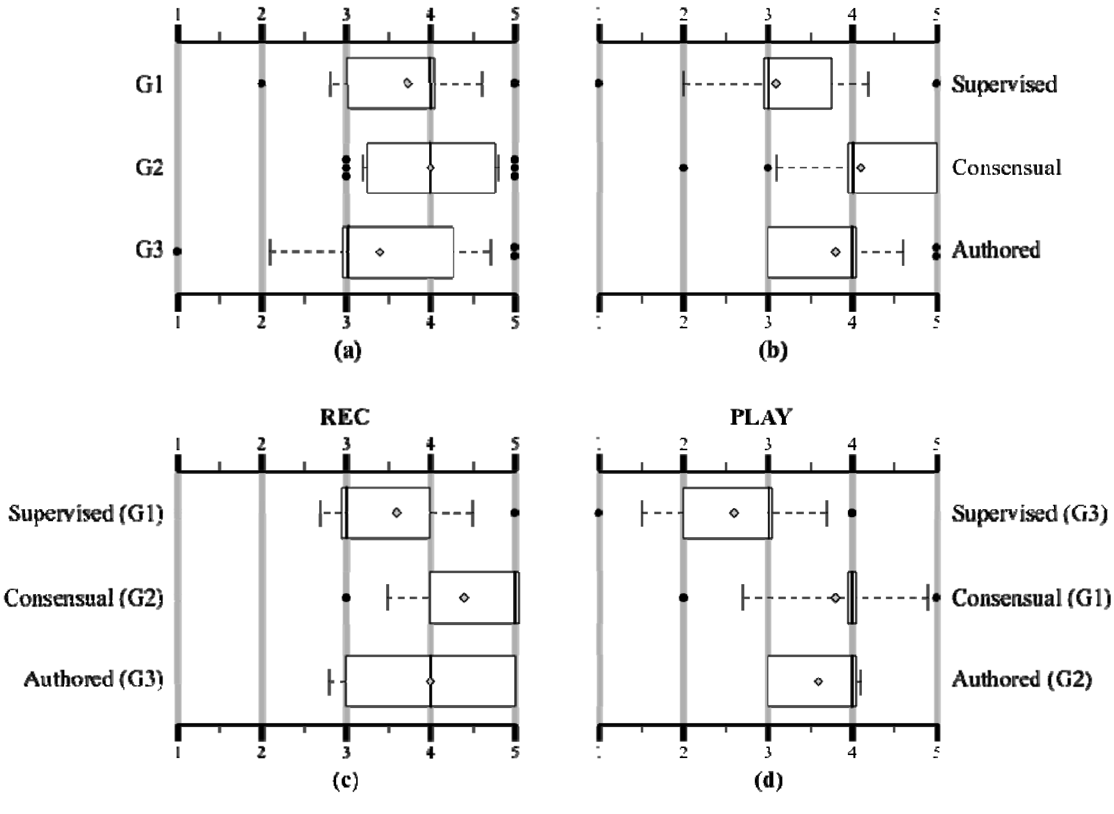

Fig. 6: How were you satisfied with the final annotations (1-5)? **(a)** shows the answers aggregated by group. **(b)** shows the answers aggregated by collaboration policy. **(c)** and **(d)** show the answers for each combination of policy and scenario

The third section of the survey concerned more specifically the collaboration policy. We asked users to assess the adequacy of the policies they had experimented with, to tell us whether they would have preferred a different one, and whether in their opinion there had been insoluble conflicts among the participants.

Fig. 7 shows, again by means of boxplots, the users' opinion on the adequacy of the different policies.

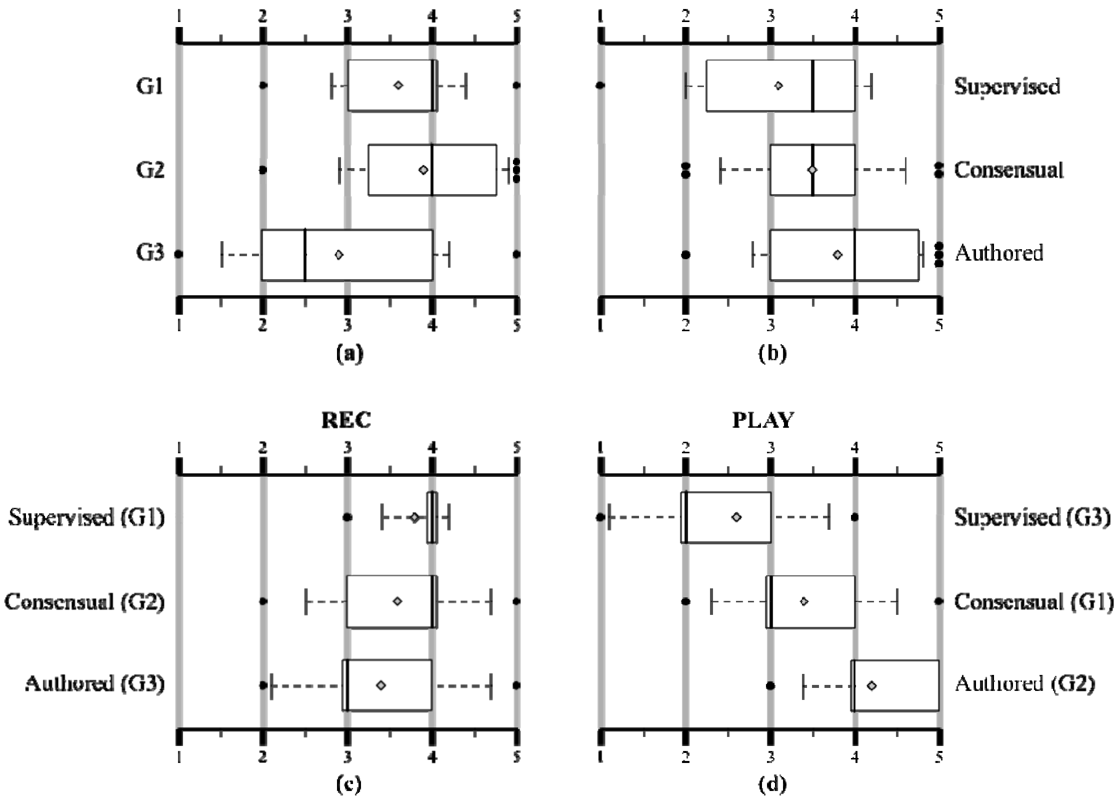

Fig. 7: Was the collaboration policy adequate (1-5)? **(a)** shows the answers aggregated by group. **(b)** shows the answers aggregated by collaboration policy. **(c)** and **(d)** show the answers for each combination of policy and scenario

Only 2 people noticed conflicts going on during the activity: 1 in group G1 in the PLAY scenario (consensual policy) and 1 in group G3 on the REC scenario (authored policy).

Only 2 subjects, among the 10 people who experimented with the **consensual policy**, felt that it was too chaotic and that a supervised policy would have been better. Both of them belonged to group G1, which used the consensual policy in the PLAY situation.

No one complained about the **authored** policy in group G2, while group G3 felt quite controversial about it: 2 people would have preferred other policies, but while one of them would choose consensus, the other would prefer the supervised one.

The **supervised** policy was found satisfactory by group G1, while 3 people in G3 did not like it:2 would have preferred the authored one, while 1 whished for a consensual policy.

Table 1 recapitulates these answers.

The last section of the survey concerned additional functionalities that, if made available (and used), may have aided the annotation process. The first two functionalities we asked about, that is communication tools and revision history, were indeed present in Google Docs, but few people thought of using them. Only 4 people in fact used communication tools (e.g. the comments function in Google Docs, or the email), and 7 explicitly said that it did not occur to them that they could (even if we briefly suggested it as a possibility in the scenario introduction). The revision history was used by 6 people; on a 5-point scale, 2 people deemed it not very useful (rating = 2), 7 people expressed an intermediate opinion (rating = 3), 3 people deemed it a useful tool (rating = 4) and 3 people rated it as very useful (rating = 5).

The other two functionalities we proposed were very difficult to simulate in Google Docs without being explicitly directed to do so (and users may have felt, if they did, that they were somehow violating the rules of the game).

The first one is a voting mechanism, where the system keeps track of votes for annotation proposals: 6 people among our subjects answered that it would have been an interesting possibility as an alternative to the consensual policy.

The second one is the possibility to keep track, along with collaborative annotations, of personal annotations, that are visible only to the user who provided them. People felt this was a better approach to avoiding conflicts or endless bickering: 11 of them in fact declared this as a desirable feature.

	G1	G2	G3
REC	Supervised 5: supervised ok	Consensual 5: consensual ok	Authored 3: authored ok 1: prefer consensual 1: prefer supervised
PLAY	Consensual 3: consensual ok 2: prefer supervised	Authored 5: authored ok	Supervised 2: supervised ok 2: prefer authored 1: prefer consensual

Table 1: Would you have preferred another policy? Which one?

4.2.3. Analysis

In this section we discuss the survey results and see how they either confirmed or shifted our preliminary ideas and assumptions, thus giving directions to our present proposal for collaborative annotation in Sem T++.

As a first remark, we can notice that the profile of participants reflected the type of users we were looking for when recruiting: all of them are reasonably familiar with the web and social media, and most of them had previous experiences with collaborative tools.

We can also notice that our subjects were significantly involved in the experience, even if we presented them with fake situations, in which they had no personal interest nor future investment. Their involvement emerges both from the average number of active interventions, and from the assessment of ease-of-use, usefulness, fun and interest of the activity itself (see Fig. 5), which on the average is either neutral or tipped toward the "positive"(right) side.

As it often happens, it is interesting to examine more closely negative remarks. Only one user found the activity useless (with a score of 2). In the open remarks, she points out that, if the goal is retrieval, she would like to search based on her own classification criteria and labels rather than on other people's. Other participants probably felt the same way, since most of them expressed interest in the possibility of maintaining personal annotations, visible only to the person providing them. Two people found the activity boring (with a score of 2); both were quite active in annotating the resources, but found the simulation too slow – one remarked that annotations converged quickly and the time we allotted for the simulation to play out was too much. In the same situation one of the participants barely had time to intervene once: this shows that people who are slower than others to enter in a process may be – willingly or inadvertently – left out by the group if the majority agrees on what has been done. This may be proper or not depending on the real-life situation; it however points to the need for better awareness mechanisms that allow to know who has done what.

Another "bored" participant remarked that other people were slow in contributing and she felt the lack of a more active coordination by the supervisor (she was playing the REC situation with the supervised policy). A possibility for this person would have been to stimulate the others by explicitly addressing them with either the chat or comment features provided by Google Docs. The participant noted that it did not occur to her that she could do so. Actually, the possibility to communicate was overlooked by most of the participants. Only 4 people chose to communicate either via personal messaging or with Google Docs comments; the others not only did not use them, but mostly remarked that they completely overlooked their existence, even when it could have helped them. Since, as one participant positively remarked, other people's annotations can help to see the resource from a different perspective, enriching each other's view on the subject matter, people should probably be actively encouraged to comment on their activity so that everyone can understand their point of view.

Four people found the activity not interesting (score 2). Two of them were the "bored" participants; the reasons were the same as mentioned before. Other two "uninterested" participants said that they were not interested in the topics presented in the scenarios. This may be due to the simulated nature of our scenario; it is however also true that in a collaborative group there may be people that are less prone to actively participate in this type of activity and prefer to leave the annotation work to others or intervene very briefly. Also, people may be forced to collaborate on topics they do not find interesting by their job or position. Uninterested, uninvolved or busy people should be able to either "pass on" the task, or receive suggestions.

As far as the collaboration policies are concerned, there are two aspects to consider: how user perceived the policy itself (Fig. 7) and how satisfied they were with the resulting annotations (Fig. 6). Notice that, dependently on the context, satisfying all the participants may or may not be a goal – if they have very different opinions, it may even be impossible. The participants themselves may accept to be (partially) unsatisfied with it (for example, because they acknowledge the right of the resource creator to have the final say on the labels) or to be annoyed at the result (for example, because they had an isolated opinion and, with a consensual policy, they felt they "lost the fight"). If we look at Fig. 7, we see that the consensual policy worked better on the REC situation, than on the PLAY situation (where 2 people explicitly remarked that they would have preferred to have a supervisor): although the interquartile distance is the same, in the REC situation the median is 4 (with all answers in the 3rd quartile equal to 4), while in the PLAY situation the median is 3 (with all answers in the 2nd quartile equal to 3). Users commented in the survey that resources in the second scenario were more difficult to annotate, due to the more elusive, "artistic" topic, and to the fact that one of the resources was a video. This apparently made it more difficult for group members to converge, although in the end they did. This may explain (see Table 1) the desire for some coordination – although, it is however interesting to note that, when the PLAY scenario was played with the supervised policy by group G3, very few were satisfied with it: since the resources could be annotated from different perspectives, the fact that the supervisor chose one perspective over the others was not received well by the participants, even those who had supported the chosen viewpoint. Also, if we look at Fig. 7 (c) and (d), none of the two groups (G1 and G2) who played with the consensual policy considered it better than the other one they experimented with (the supervised policy for group G1 and the authored one for group G2). However, maybe not surprisingly, it gets better results with respect to the other two policies when it comes to satisfaction in the resulting annotations, as people felt their opinion was represented in the result (see Fig. 6 (b)).

The supervised policy posed no problem to users in G1, on the REC situation, while 2 people complained about it in G3, on the PLAY situation, where (see Fig. 7 (d)) we see the worst ratings for adequacy – actually the only one with a median lower than 3.
Also in this case the more elusive and controversial nature of the PLAY resources played a role: as we remarked above, when playing the supervisor we chose one of the many perspectives from which the resource could be annotated, and people did not like that the "discussion" they had carried out by adding and editing annotations, trying to bring in as many perspectives as they could, was disregarded and cut off by the supervisor's decision. By observing the revision history, we had the impression that users were trying to solve conflicts in the PLAY scenario by adding as many *objects of discourse* as they could and by using a very long string of text as *topic*; in other words they were bringing everything in. As supervisors, we could let this be (but this would be similar to a consensual policy, and we have seen that for group G1 this did not work either) or make a choice. The fact that people were unhappy in this case and perfectly happy in the REC situation, tells us that the perceived adequacy of the supervised policy depends more on the attitude of the supervisor than on the policy itself. This is easily seen in Fig. 7 (c) and (d): in the REC scenario, the supervised policy is consistently perceived as adequate (actually more than the other two), while in the PLAY scenario the supervised policy is mostly perceived as inadequate, and definitely less adequate than the other two. This is reflected also on the degree of satisfaction with the final labels (see Fig. 6 (c) and (d)).
The authored policy was perceived as adequate in the PLAY situation (which was not easy), obtaining the highest ratings (Fig. 7 (d)). It was more controversial in the REC situation (Fig. 7 (c)), played by group G3, where it gets the lowest ratings with respect to the other two, although in this scenario there is less distance among the adequacy ratings of the different policies. On the overall, in fact, the authored policy is perceived as the most adequate (Fig. 7 (b)), leading on the overall to satisfying annotations (Fig. 6 (b)), even if not as satisfying as with the consensual policy.
It is important to point out that our goal was not to choose the best policy, but to assess weaknesses and strengths of each of them, as well as to understand whether we needed to discard some, propose other ones, or provide support to improve a certain collaboration mode.
Both the consensual and the supervised policy showed their weaknesses in controversial situations: from the participants' remarks, we gathered that both scenarios could have been improved by increased awareness and communication. Increased awareness concerning the contribution of each participant to the current version (that is, knowing which annotations were expressed by whom) helps in recognizing the different perspectives, and provides motivation for acknowledging them as much as possible. A greater degree of communication, both within the group and with the

supervisor (in the supervised case), can help to clarify the reasons behind different perspectives, finding common points and compromises that do not leave out potentially useful points of view. Interestingly enough, these two points were already raised while discussing the quality of the collaborative experience. The choice to communicate or not remains in the hands of the participants, but it is clear that providing encouragement to do so can improve the level of interaction among people. Providing awareness about each person's contribution can be more challenging – as we already pointed out, in a sense this information is contained in the revision history, but these tools, typically, do not help in recognizing the perspective that each person brought to the annotations. It may take some time to realize, for example, that a participant constantly tried, with each intervention, to bring in a certain point of view that others constantly discarded. The need for establishing and clarifying individual perspectives brings in also the issue of "personal annotations": most of our users saw this as a useful feature, mainly because it would allow them to use their personal keywords to retrieve resources, rather than the collective ones. However one could imagine to blend the border between collective and personal, where personal annotations – if made visible to others – can represent individual perspectives, and can help in building a collective version that partly represents them and partly enriches them. Such an approach would also subsume the matter of "votes": most of our subjects said they would not want a voting mechanism. Those few who liked the idea of voting reported that they felt the need to understand, in controversial cases, how many people were in favor of a given point of view. Also, they felt that this information could not only shift the consensual policy toward a more democratic approach, but it could also improve the supervised and authored policies, because it would allow the coordinator (supervisor or creator) to see how many people liked or disliked a certain annotation. However, most participants thought that a voting mechanism would be too cumbersome for a small group of people, and that fostering a discussion that makes room for different perspectives and for creative inclusive solutions is the ideal approach.

4.3. Collaborative Semantic Annotation: Guidelines and Implementation in Sem T++

4.3.1. Guidelines

On the basis of the analysis of the user study results presented in section 4.2.3, we defined the following guidelines. They guided us in the design and implementation of the collaborative semantic annotation support tool.

According to our user study, such systems should:

1. Provide users with a **choice** within at least the following **collaboration policies**:

- Consensual, where the possibility of editing the semantic description of table objects is never closed (or it can be considered closed only when consensus is reached among all participants).
- Authored, where the possibility of editing the semantic description of table objects is closed by the object creator, who is the only one who can open it again. In this case, for each table object a possibly different person is responsible for the collaborative annotation process.
- Supervised, where the possibility of editing the semantic description of table objects is closed by the table supervisor, who is the only one who can open it again. In this case, the same person is responsible for the collaborative annotation process of all table objects.

This guideline emerges from the observation that, although certain policies lead to greater user satisfaction, no policy was actually perceived as intrinsically negative or unproductive by our test subjects. Complaints were mainly raised not on the policy itself, but rather on the lack of communication and in the difficulty of understanding what was happening on the table.

2. Provide users with the possibility of an **overall approval** of the annotations concerning a resource, meaning something like "I was here, I saw what you did and it is ok for me". This solves the problem of users waiting indefinitely for someone's opinion when the person simply has nothing to object on what has been done. Therefore it increases awareness concerning the collaboration and speeds up the process.
3. Provide users with a feature enabling them to **see "at a glance" each participant's responsibility** in annotations. This was one of the major points pointed out by our test subjects: this feature, in fact, together with the previous one, significantly raises the awareness level about each participant's contribution and perspective, showing who has contributed (or, at least, approved) and who has not, without having to sort out each editing action in a full revision history.
4. **Facilitate and encourage communication** among participants. A major point raised by our user study was that users would have used some communication tool if only they had remembered they could do it. Thus the system should make communication tools easily available and should encourage users to use them, possibly actively prompting them to provide comments, explanations, opinions when they edit annotations.

5. Provide users with the possibility of saving **private annotations**, representing their personal view over resources. Most of the participants in our user study, in fact, were interested in handling annotations visible only to the person who provided them.

4.3.2. Implementation

On the basis of the guidelines presented above, in Sem T++ we decided to provide all of the **three policies** that were experimented within the user study (guideline 1), namely: *consensual*, *authored*, *supervised*.

We decided not to include more complex policies, such as voting mechanisms, delegation, or workflow management; such mechanisms will be possibly taken into account if additional scenarios will require them.

The choice of the collaboration policy can be configured by the table creator through the **configuration panel** (currently a very simple web form); if the supervised policy is selected, a table supervisor must be nominated among table participants. Such configuration can be modified at any time during the table lifecycle.

Once the policy has been selected, table participants can annotate resources by adding or deleting property values. Fig. 8 shows the window enabling users to edit the properties of a table object (*Recycling Tips*); in particular the panel for editing objects of discourse (i.e., "what the resource talks about") is displayed. At the bottom of the window, a checkbox enables users to mark the overall resource annotation as **approved** (guideline 2). This feature flanks candidate suggestions (see Fig. 4) in alleviating the repetitiveness of the annotation task.

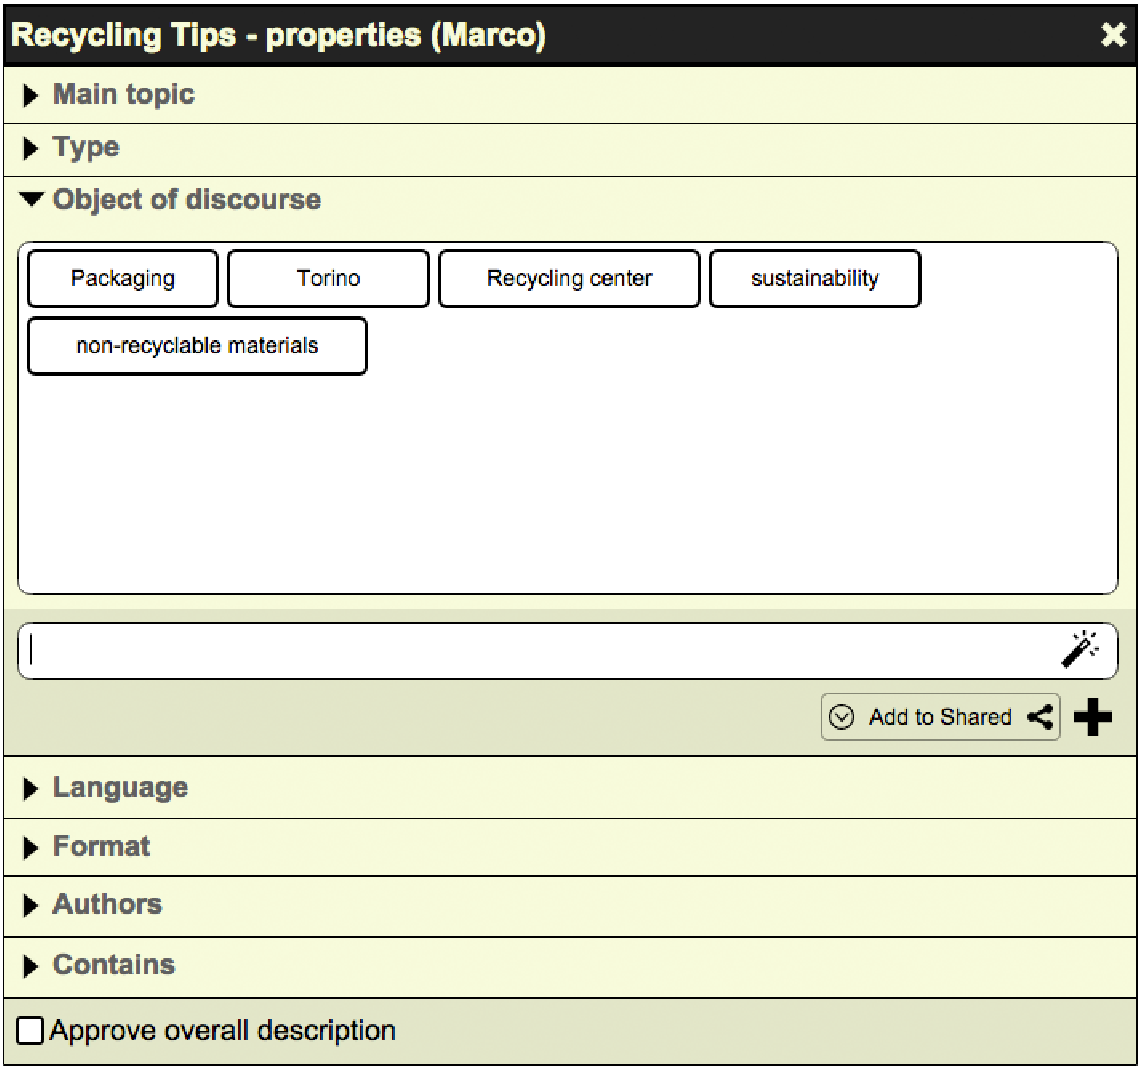

Fig. 8: Sem T++ user interface: the window enabling users to view/edit resource properties

Fig. 8 refers to a table where collaborative annotations are handled by a *consensual* policy. However, if a table has been configured with a *supervised* or *authored* policy, the supervisor or the author can "freeze" the annotation of a resource (by clicking on a "close" button), thus making all property values read-only.

Moreover, Sem T++ provides a standard **revision history** for each table object, showing the annotation activity performed by table participants, and also a ***history at-a-glance***, aimed at showing the author of each annotation (guideline 3). This simplified view of the change history concerning semantic descriptions provides a first-glance view of the "responsibilities" concerning property values. In the current prototype, the *history at-a-glance* is implemented in a straightforward way: in the panel enabling users to see and modify (when in editable status) property values (see Fig. 8), right clicking the value shows a tooltip with the author of the annotation.

As we mentioned in Section 3.1, Sem T++ provides different standard **awareness tools**, ranging from icon highlighting to notification messages. When something occurs in the semantic description of a table object (e.g., a value has been changed, editing has been closed or re-opened), table participants are notified. This is very important in order to guarantee that supervised collaboration policies, as well as the semantic description approval mechanism, work properly.

Moreover, Sem T++ provides, on each table, a range of **tools supporting communication** among table participants, i.e., a Blackboard for posting asynchronous messages, a table Chat, and free-text Comments which can be attached to table objects (see Section 3.1). All these tools can be exploited to facilitate the discussion among users. In particular, in Sem T++, when a user modifies a value in the semantic representation of a table object, the table invites her to write a comment in order to explain it, by a pop-up window linked to the Comments tool (guideline 4).

The present design and implementation of Sem T++ acknowledges most of the guidelines emerged from the qualitative user study. There is however a significant point that we did not tackle, namely the possibility for users to keep personal annotations (guideline 5). We actually believe that this issue is so relevant that it requires further study, thus we will consider it in our future work.

5. Conclusions

In this paper we faced the issue of collaborative semantic annotation of shared resources, taking place within small teams of people cooperating through the Sem T++ environment. Semantic annotation in Sem T++ is driven by a formal ontology, defining the properties of information objects. While the values of some properties can be automatically inferred, properties that describe the content of the resource need to be agreed upon by the team members. This led us to study the issue of collaborative annotation. In the paper we first presented a qualitative user study where we observed the process of collaborative annotation in small groups of people, under different assumptions, especially for what concerns the agreement policy. Our subjects experimented with collaborative annotation with a consensual policy (where no one establishes a "final" annotation), a supervised policy and an authored policy (where, in both cases, someone closes the process establishing a final annotation; in the former case an external supervisor, in the latter the resource creator). The results of the user study allowed us to draw a few guidelines on which aspects of this process can be supported or improved by design choices when implementing collaborative annotation. We can summarize them as follows:

- it is important that team members are aware of what has been done and who has done what, as well as who agrees with the present annotation and who does not;

- team members, and the supervisor if there is one, should be encouraged to communicate the perspective they are trying to express with their annotations;
- it should be possible, for each team member, to keep her own perspective (as a personal annotation) along with the shared one.

In the paper we present a collaborative tool for annotation that implements most of these features. In particular, it implements all of the three policies described above, giving users the possibility to choose one in the table configuration. It supports the collaborative process by improving awareness of other's contributions and fostering communication. The tool also support users with automated suggestions of possible annotations, and by allowing participants to simply approve what others did, without having to explicitly provide their contribution.

As already partially mentioned in previous sections, some important issues, raised by the user study, are still to be investigated in detail, and represent future development of our system. Among these, the most important is the possibility of maintaining personal semantic descriptions of shared resources, representing personal perspectives. We are currently investigating formal semantic mechanisms that can handle the overlay of personal and shared annotations.

Moreover, many of the design decisions concerning collaborative semantic annotation described in this paper pose great challenges concerning user interaction: we are thus studying the impact of the implemented mechanisms on user interaction, also taking into consideration multi-device user interfaces, in order to reach a better user experience.

References

[1] W. Wahlster, A. Dengel, Web 3.0: Convergence of Web 2.0 and the Semantic Web, Technology Radar II (2006) 1–23.

[2] J. G. Breslin, A. Passant, S. Decker, The Social Semantic Web. Springer, Heidelberg, 2009.

[3] A. Goy, G. Petrone, M. Segnan, Oh, no! Not Another Web-based Desktop!, in: Proceedings of the 2nd Int. Conf. on Advanced Collaborative Networks, Systems and Applications, XPS Press, Wilmington, DE, 2012, pp. 29–34.

[4] V. Kaptelinin, M. Czerwinski (Eds.), Beyond the Desktop Metaphor, MIT Press, Cambridge, MA, 2007.

[5] D. R. Karger, Haystack: Per-User Information Environments Based on Semistructured Data, in V. Kaptelinin, M. Czerwinski, M. (Eds.), Beyond the Desktop Metaphor, MIT Press, Cambridge, MA, 2007, pp. 49–100.

[6] J. E. Bardram, From Desktop Task Management to Ubiquitous Activity-Based Computing, in V. Kaptelinin, M. Czerwinski, M. (Eds.), Beyond the Desktop Metaphor, MIT Press, Cambridge, MA, 2007, pp. 223–260.

[7] S. Voida, E. D. Mynatt, W. K. Edwards, Re-framing the Desktop Interface Around the Activities of Knowledge Work, in Proceesings of the 21st ACM Symposium on User Interface Software and Technology, ACM Press, New York, NY, 2008, pp. 211–220.

[8] A. Goy, G. Petrone, M. Segnan, A Cloud-Based Environment for Collaborative Resources Management. Int. J. Cloud Applications and Computing 4(4), 2014, 7–31.

[9] L. Sauermann, A. Bernardi, A. Dengel, Overview and Outlook on the Semantic Desktop, in Proceedings of the 1st Workshop on The Semantic Desktop at ISWC 2005, vol. 175, CEUR-WS, 2005.

[10] L. Drăgan, R. Delbru, T. Groza, S. Handschuh, S. Decker, Linking Semantic Desktop Data to the Web of Data, in L. Aroyo, C. Welty, H. Alani, J. Taylor, A. Bernstein, L. Kagal, N. Noy, E. Blomqvist (Eds.), The Semantic Web – ISWC 2011. LNCS, vol. 7032, Springer, Heidelberg, 2011, pp 33–48.

[11] R. Koning, P. Grosso, C. de Laat, Using ontologies for resource description in the CineGrid Exchange, Future Generation Computer Systems, 27(7), 2011, 960–965.

[12] Y. F. Li, G. Kennedy, F. Ngoranb, P. Wud, J. Hunter, An ontology-centric architecture for extensible scientific data management systems, Future Generation Computer Systems 29(2), 2013, 641–653.

[13] M. Uflacker, A. Zeier, A semantic network approach to analyzing virtual team interactions in the early stages of conceptual design, Future Generation Computer Systems 27(1), 2011, 88-99.

[14] H. Cunningham, D. Maynard, K. Bontcheva, V. Tablan, N. Aswani, I. Roberts, G. Gorrell, A. Funk, A. Roberts, D. Damljanovic, T. Heitz, M. A. Greenwood, H. Saggion, J. Petrak, Y. Li, W. Peters, Text Processing with GATE (Version 6), 2011 (gate.ac.uk).

[15] P. Fragkou, G. Petasis, A. Theodorakos, V. Karkaletsis, C. Spyropoulos, Boemie ontology-based text annotation tool, in N. Calzolari, K. Choukri, B. Maegaard, J. Mariani, J. Odjik, S. Piperidis, D. Tapias (Eds.), Proceedings of the Int. Conf on Language Resources and Evaluation (LREC'08), European Language Resources Association (ELRA), 2008.

[16] K. Bontcheva, H. Cunningham, I. Roberts, V. Tablan, Web-based collaborative corpus annotation: Requirements and a framework implementation, in R. Witte, H. Cunningham, J. Patrick, E. Beisswanger, E. Buyko, U. Hahn, K. Verspoor, A. R. Coden (Eds.), New Challenges for NLP Frameworks, 2010.

[17] G. Petasis, The SYNC3 Collaborative Annotation Tool, in Proceedings of the Int. Conf. on Language Resources and Evaluation (LREC'12), European Language Resources Association (ELRA), 2012, 363–370

[18] J. Chamberlain, M. Poesio, U. Kruschwitz, Phrase detectives - a web-based collaborative annotation game, in Proceedings of I-Semantics, 2008.

[19] V. Uren, P. Cimiano, J. Iria, S. Handschuh, M. Vargas-Vera, E. Motta, F. Ciravegna, Semantic annotation for knowledge management: Requirements and a survey of the state of the art, J. of Web Semantics, 4(1), 2006, 14–28.

[20] O. Corcho, Ontology based document annotation: trends and open research problems, Int. J. of Metadata, Semantics and Ontologies,1(1), 2006, 47–57.

[21] M. E.I. Kipp, D. G. Campbell, Patterns and Inconsistencies in Collaborative Tagging Systems: An Examination of Tagging Practices, in Proceedings of the American Society for Information Science and Technology, 43(1), 2006, 1–18.

[22] R. Pak, S. Pautz, R. Iden, Information organization and retrieval: An assessment of taxonomical and tagging systems, Cognitive Technology 12(1) (2007), 31–44.

[23] A. Civan, W. Jones, P. Klasnja, H. Bruce, Better to organize personal information by folders or by tags? The devil is in the details, American Society for Information Science and Technology 45, 2009, 1–13.

[24] S. A. Golder, B. A. Huberman, Usage patterns of collaborative tagging systems, J. of Information Science, 32(2), 2006, 198–208.

[25] H. Kim, J. G. Breslin, S. Decker, J. Choi, H. Kim, Personal Knowledge Management for knowledge workers using social semantic technologies, Int. J. of Intelligent Information and Database Systems 3(1), 2009, 28–43.

[26] F. Abel, N. Henze, D. Krause, M. Kriesell, Semantic Enhancement of Social Tagging Systems, in V. Devedžić, D. Gašević (Eds.), Web 2.0 & Semantic Web, Springer, Heidelberg, 2010, pp. 25–56.

[27] A. Y. S. Su, S. J. H. Yang, W. Y. Hwang, J. Zhang, A Web 2.0-based collaborative annotation system for enhancing knowledge sharing in collaborative learning environments, Computers & Education 55, 2010, 752–766.

[28] M. Buffa, F. Gandon, G. Ereteo, P. Sander, C. Faron, SweetWiki: A semantic wiki, Web Semantics, 6(1), January, 2008, 84–97.

[29] F. Papazian, R. Bossy, C. Nédellec, AlvisAE: a collaborative web text annotation editor for knowledge acquisition, in Proceedings of the 6th Linguistic Annotation Workshop, 2012, 149–152.

[30] G. Petasis, V. Karkaletsis, G. Paliouras, I. Androutsopoulos, C. D. Spyropoulos, Ellogon: A New Text Engineering Platform, in Proceedings of the Int. Conf. on Language Resources and Evaluation (LREC'02), European Language Resources Association (ELRA), 2002, 72–78.

[31] L. Ardissono, G. Bosio, A. Goy, G. Petrone, Context-Aware Notification Management in an Integrated Collaborative Environment, in Proceedings of the workshop on Adaptation and Personalization for Web2.0, CEUR-WS, 2010, 23–39.

[32] M. Reynolds, How to Collaborate With Google Docs, 2013, SpinWeb Blog, http://blog.spinweb.net/how-to-collaborate-with-google-docs, accessed 30/11/2014.

[33] A. Goy, D. Magro, G. Petrone, M. Segnan, Semantic representation of information objects for digital resources management, Intelligenza Artificiale, 8(2), 2014, 145–161.

[34] S. Borgo, C. Masolo, Foundational choices in DOLCE, in S. Staab, R. Studer (Eds.), Handbook on Ontologies – Second Edition, Springer, Heidelberg, 2009, pp. 361–381.

[35] A. Gangemi, S. Borgo, C. Catenacci and J. Lehmann, Task Taxonomies for Knowledge Content, Metokis Deliverable D07, 2005.

[36] D. Magro and A. Goy, A Core Reference Ontology for the Customer Relationship Domain, Applied Ontology, **7**(1), 2012, 1–48.