

Tags4Tags: Using Tagging to Consolidate Tags

Leyla Jael Garcia-Castro¹, Martin Hepp¹, Alexander Garcia²

¹ E-Business and Web Science Research Group, Universität der Bundeswehr München,
D-85579 Neubiberg, Germany
leyla.garcia@unibw.de, mhepp@computer.org

² Department of Computational Linguistics, University of Bremen
D-28359 Bremen, Germany
cagarcia@uni-bremen.de

Abstract. Tagging has become increasingly popular and useful across various social networks and applications. It allows users to classify and organize resources for improving the retrieval performance over those tagged resources. Within social networks, tags can also facilitate the interaction between members of the community, *e.g.* because similar tags may represent similar interests. Although obviously useful for straightforward retrieval tasks, the current meta-data model underlying typical tagging systems does not fully exploit the potential of the social process of finding, establishing, challenging, and promoting symbols, *i.e.* tags. For instance, the social process is not used for establishing an explicit hierarchy of tags or for the collective detection of equivalencies, synonyms, morphological variants, and other useful relationships across tags. This limitation is due to the constraints of the typical meta-model of tagging, in which the subject must be a Web resource, the relationship type is always *hasTag*, and the object must be a tag as a literal. In this paper, we propose a simple yet effective extension for the current meta-model of tagging systems in order to exploit the potential of collective tagging for the emergence of richer semantic structures, in particular for capturing semantic relationships between tags. Our approach expands the range of the object of tagging from Web resources only to the union of (1) Web resources and (2) pairs of tags, *i.e.*, users can now use arbitrary tags for expressing typed relationships between a pair of tags. This allows the user community to establish similarity relations and other types of relationships between tags. We present a first prototype and the results from an evaluation in a small controlled setting.

Keywords: Social Web, folksonomy, tagging, meta-model, emergent semantics, conceptual graphs, Semantic Web, Web 2.0.

1 Introduction

Nowadays social tagging systems (STS), and the resulting knowledge structures known as folksonomies [1], are widely used on the Web. Tagging typically works by assigning short lexical elements to resources in a collaborative environment, mainly for document retrieval. Popular sites focus on tagging Web resources (*e.g.* Delicious, <http://www.delicious.com/>, and Connotea, <http://www.connotea.org/>), images (*e.g.* Flickr, <http://www.flickr.com/>), or blogs and other user-generated content (*e.g.* Technorati, <http://www.technorati.com/>). Recently, respective technology has also

been used in corporate networks such as the Electricité de France Intranet [2], where tags were used in blogs to promote knowledge sharing inside the organization.

It can be assumed that the popularity of tagging is not only due to the simplicity of the tagging operation itself, but also because tags effectively facilitate search and navigation over tagged resources [3]. From the technical perspective, there are several attractive features of tagging systems that create added value for users: First, the use of URIs for resources provides reliable, unique identifiers for documents, which allows for the consolidation of meta-data. Second, the sites provide a collaborative environment with an explicit representation of users, which allows discovering implicit relationships, *e.g.* networks of users with similar skills, tasks, or interests. Third, tagging systems provide simple yet effective support for the emergence of consensus on (i) the exact lexical form of a tag and (ii) the appropriateness of a tag for a certain resource based on collaborative filtering and recommendations. This helps to avoid orphaned tags and reduces lexical or morphological variations; at the same time, it keeps up with the high agility and the good coverage of rare but still relevant elements, *i.e.* such that are on the long tail. Centralized approaches, including classical ontology-based solutions often lag behind in their coverage of novel or specific domain elements [4]. Furthermore, as tags work like bookmarks or indexes, they help to reduce spam-induced noise in search engines and enable text-based queries over elements like images [5]. Moreover, tagging does not impose rigid categories or a controlled vocabulary on users but gives to users the possibility to freely create and associate terms, *i.e.*, descriptors, to resources.

Although tagging has proven to provide significant benefits, there are also relevant limitations of the current state of technology. Typical problems are (i) tag ambiguity, (ii) missing links between multiple synonyms, spelling variants, or morphological variants, and (iii) variation in the level of granularity and specificity of the tags used caused by differences in the domain expertise of agents [2, 3, 6, 7]. These limitations are problematic for *e.g.* (i) developing intelligent user interfaces for annotations, (ii) improving navigation and querying based on annotations, and (iii) integrating content from diverse and heterogeneous data sources [6].

Additional formal structures may help to overcome some of problems mentioned above [5, 8, 9]. A main question, however, is whether such formal structures are imposed explicitly in the tagging stage or derived implicitly by mining tagging data.

While there exist many proposals for the latter approach, we propose to expand the underlying meta-model of tagging systems from attaching tags to resources only to attaching tags to resources and *arbitrary pairs of tags*, *i.e.*, pairs of the form (tag, tag). Our motivation is to exploit the positive technical and social effects of tagging for the construction and the management of more powerful conceptual structures in information systems and on the Web.

Some of the expected advantages of this model are: (i) supporting the emergence of explicit relationships between tags, (ii) adding meaning to numerical tags, (iii) building a “tagsonomy”, *i.e.*, a conceptual graph of tags, including relationships between them, and (iv) improving the basis for adding formal semantics to tags by mining techniques.

We expect this to improve the retrieval performance of tagging systems and to help building conceptual graphs from a set of tags. This may turn STS into true sources of collective intelligence. Such likely requires the aggregation and recombination of data

collected from annotations in order to create new knowledge and new ways of learning that individual humans cannot achieve by themselves [10].

This paper is organized as follows. Section 2 introduces our model and motivates our approach. Section 3 summarizes our preliminary evaluation by means of a controlled experiment on establishing relations between tags. Section 4 summarizes related work and discusses our contribution. Section 5 sketches future directions for research.

2 A Vision to Expand the Scope of Taggable Objects

Currently, STS allow agents (A), i.e. users, to add tags (T) to resources (R); each respective activity is called tagging (TA). This simple setting already allows a high degree of variation as presented in Fig. 1.

Some of the strengths of STS arise from those combinations such as promoting serendipity, facilitating convergence, and supporting collaboration by means of filters and recommendations based on existing tags. Since users can share their tags with the community, they are building not only a knowledge representation for themselves but are also helping others to discover associations that were not previously known. This increases the collectively available expertise [11] and supports the social reinforcement by means of enabling social connections based on common interests [8]. In other words, the aggregation of many individual user contributions can by itself create an added value in STS [10].




Fig. 1. Variations in tagging resources.

A STS can be represented as a graph where agents (A), resources (R), and tags (T) are the vertices (V) and tagging activities (TA) are the edges. The predominant tagging meta-model relies mainly on a triple of (i) agents, i.e. users, (ii) resources, and (iii) tags. It is also possible to include the system where the tagging took place and a polarity to assign negative or positive values to tags as Gruber proposes [5]. Some authors add relations between tags such as *relatedTo* and *equivalentTo* which are taken from Newman's model [12]. The typical meta-model used today can be formalized as follows:

$$\text{STS} = \langle V, \text{TA} \rangle \mid V = A \cup T \cup R \text{ and } \text{TA} \subset \{(A, T, R)\} \quad (1)$$

Although tagging has proven to provide significant benefits, there are also important limitations. Several of these limitations are caused by the constraints in the

current meta-model, as the subject of the tagging must be a Web resource, the relationship type is always *hasTag*, and the object must be a tag.

Often, social tagging suffers from a lack of structure and contextualization. For example, we often do not know for which purpose a particular user attached a certain tag to a certain resource. The assignment may *e.g.* be relevant or valid only in the context of a particular task. Since all tags are organized in a shallow way, the navigation, querying, and retrieval of resources is limited [2, 6] and free relations between tags cannot be established by the same social process. Also, it has been observed that people need to contextualize communication with other people because that fosters the creation of new knowledge [10]. Being able to tag other agents' previous tagging activities would support that.

The key motivation for our approach is to keep the simplicity and popularity of tagging while using them on richer conceptual structures, instead of solely trying to derive those structures from tagging data by mining techniques. In our opinion, the current meta-model does not fully exploit the potential of tagging for finding, establishing, challenging and promoting symbols in a community. For instance, the social convergence in STS is not used for establishing an explicit hierarchy of tags or for the collective detection of equivalencies, synonyms, morphological variants, and other useful relationships across tags.

We propose to expand the current meta-model by opening up the range of taggable objects *from resources only to resources and relations between tags*. We aim at relating entities from this expanded set in a semantic manner by means of that approach. Our model is shown in Fig 2.

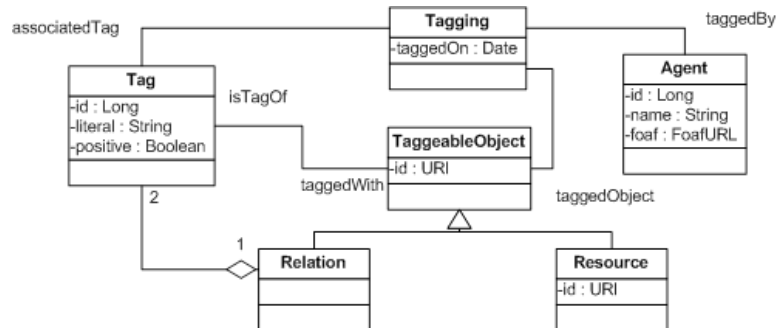


Fig. 2. Expanding the scope of taggable objects.

A STS to support our model can be still represented as a graph. As before, agents (A), Web resources (R), and tags (T) are the vertices (V) and tagging activities (TA) are the edges. One new vertex is required in order to represent relations between tags (RT). It can be formalized as follows:

$$STS = \langle V, TA \rangle \mid V = A \cup T \cup R \cup RT \text{ and } RT \subset \{T \times T\} \quad (2)$$

Our model intends to widen the scope of social tagging. The subject of the tagging remains an agent and the predicates remains a tag, but the object now can be either a Web resource or a type of relation between tags. The relationship type can still be

hasTag but it is also possible to add new relation types between any pair of tags. This minimal change should facilitate the usage of our new model by users familiar with traditional tagging. A summary of possible tagging scenarios with our extended model is shown in Table 1.

Table 1. Summary of possible tagging activities with Tags4Tags.

Subject	Predicate	Object	Example
Agent	Tag	Resource	(agent1, travel, http://vacations.com)
		(Tag, Tag)	(agent1, englishToSpanish, (tag:travel, tag:viaje))

Since our model is based on the current meta-model, it will likely be possible to adapt and apply (i) existing approaches to derive formal structures from tagging data such as FLOR [6] and SCARLET [13], (ii) normalization and disambiguation techniques such as [14-16], (iii) the addition of meaning to tags by using URIs [9], and (iv) techniques and tools for the tag data consolidation among platforms [8].

With the proposed extension, we basically allow people to build and maintain conceptual graphs based on tagging and complement this by social mechanisms for convergence. The resulting networks can be an important starting point to allow better retrieval and more sophisticated processing, and will likely allow more powerful approaches for deriving formal structures.

3 Prototype and Implementation

In order to evaluate our model, we developed a first prototype with the main goal of analyzing how well people are able to relate objects by tags representing the type of relationship. The model was initially populated with resources, agents, and tags related to “travel”, which we collected via the Connotea API (<http://www.connotea.org/wiki/WebAPI>), and relationships which we collected from participants by means of a Java Web-based application. The architecture of our prototype is presented in Fig. 3.

In order to capture relationships between taggable objects, we provide two columns: The left-hand side corresponds to the subject of the relationship and right-hand side corresponds to the object. The prototype was presented as a game where participants had to find as many relations as they could in a given period of time. The tool provides a set of recommended tags for likely relationships as well as suggestions based on existing tags in the system, see Fig. 4. Note that the predefined relations like *isPartOf* are also just tags. A formal meaning can be associated to those based on the outcome of the social tagging process.

The prototype was implemented as a Web-based application using Java 1.6 as the development language, the Spring Framework (<http://www.springsource.org/>) and Velocity 1.4 (<http://velocity.apache.org/>) to manage the Model-View-Controller architecture; ExtJS (<http://extjs.com/>), BoxOver (<http://boxover.swazz.org/example.html>), and Autosuggest BSN (http://www.brandspankingnew.net/specials/ajax_autosuggest/ajax_autosuggest_autocomplete.html) libraries for the user

interface, Direct Web Remoting (<http://directwebremoting.org/>) for AJAX, and Jena (<http://jena.sourceforge.net/>) as the underlying Semantic Web framework.




Fig. 3. Prototype architecture.

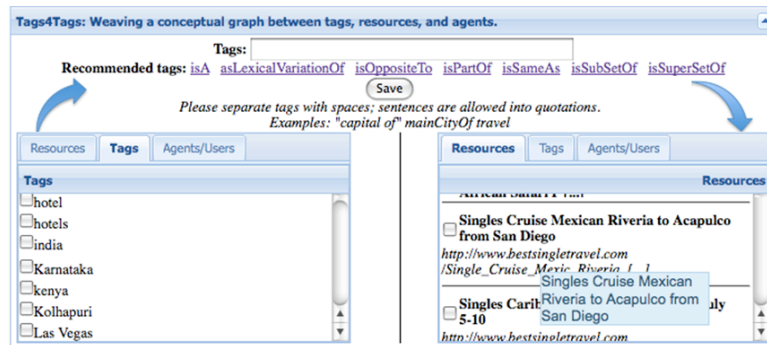


Fig. 4. A screenshot of the prototype to relate taggable objects.

4 Evaluation

The prototype described in the previous section was used to test whether the T(T, T) pattern, *i.e.* tags attached to pairs of tags, can be used to consolidate tag sets and to elicit useful relationships between tags. Those could be used for more powerful tag-based retrieval. Thus the main goal was to find out whether our approach is a feasible way of gathering user input to capture equivalencies and relationships, such as narrower and broader relations. We wanted to evaluate whether (i) average computer users are able to grasp our idea and employ it with minimal instructions, and whether (ii) the collected data is of sufficient quality to be useful.

4.1 Methodology

We recruited ten individuals from our university, both employees and students, and asked them to spot and enter as many relations as they could in a predefined period of time. All participants, seven bachelor students, one PhD student, and two researchers, had experience using the Web and some of them using tagging systems.

A set of 92 tags related to “travel” was randomly selected via the Connotea API as the initial data. We only considered tags longer than four characters and resources

with titles longer than five characters. Once our model was populated with the data, we developed the Web-based prototype to allow people to relate tags easily. The prototype, already presented in the previous section, was build with Eclipse (<http://www.eclipse.org>) and was tested manually.

Participants were asked independently to take part in the experiment and they received brief oral instructions only. The experiment was conducted in a sequential order on a single machine. All data was collected in an RDF file, and relations established by the participants were automatically loaded as part of the initial data for the following participants. This happens because relations are established by means of tags, thus they become part of the set of tags as well.

4.2 Results

All participants used the application with ease and established multiple different relations between tags. We observed that the task was harder for the first participants since they had only the initial data, and it got increasingly easier for the later participants since they could reuse relationships previously collected.

The relations between similar tags were quite consistent among users instead of the use of different tags, see Fig. 5. For instance, people attached similar tags such as “canbereachedby” and “canbevisitedby” to relate places like Berlin and Antigua to the tag “airline”. Another example is the relation between “vacations” and “vacation” with tags such as “singular” and “isLexicalVariationOf”. Someone with more experience in triples construction could have added a tag “isSubSetOf” between “singular” and “isLexicalVariationOf”, *i.e.* (“isSubSetOf”, (singular, isLexicalVariationOf)), in order to consolidate data.



Fig. 5. Some of the results found.

In order to achieve a better use of some of the recommended relationships such as “isA” and “isPartOf” as well as others proposed by users themselves, it could be necessary to identify the domain or context of taggable objects. It could also be useful to understand the meaning of some tags such as “avianFlu”. This particular tag was related to the tag “airline” with the tag “isA”, *i.e.* (“isA”, (avianFlu, airline)), by one of the participants; however this tag was initially attached to a resource related to a virus known as “avian influenza”. Additionally, background and education are also important to understand some tags; a biologist or a medical doctor would hardly have misunderstood the meaning of “avianFlu”. We assume that in Web-wide tagging systems based on our approach, the increased mass of tagging data will simplify filtering out noise and contradictions more easily.

5 Discussion and Related Work

5.1 Related Work

There are different approaches to improve social tagging by means of structured meta-models and ontologies. Approaches in this vein can be classified into five main groups: (i) modeling tagging [12], (ii) augmenting the user-contributed data [5], (iii) adding meaning to tags [9, 17], (iv) adding meta-data in order to improve retrieval, information exchange and knowledge reuse [18], and (v) enhancing sharing and reuse of social tagging data through different platforms [8].

Newman [12] proposes an ontology to “model the relationship between an agent, an arbitrary resource, and one or more tags”. His model represents tags, agents, resources and tagging activities as classes and relates them via object properties; some relations between tags such as *relatedTo* and *equivalentTo* are also modeled. A tagging activity is defined as a triple, which corresponds to a resource tagged by an agent with an associated tag. The conceptual approach of Newman’s ontology is based on the theoretical work by Gruber [5] and is taken as the baseline for other models because of its simple but comprehensive nature. Gruber’s approach [5] is broader than Newman’s because it includes other elements rather than tags, agents, and resources. Hence Gruber works with a quintuple; this quintuple incorporates information about (i) the system where the tagging took place and (ii) a polarity to represent positive and negative tags. Gruber’s ontology is part of the TagCommon project (<http://tagcommons.org/>).

MOAT [9] is one of the models that extends Newman’s ontology. Its aim is to semantically enrich content from free tagging by means of providing a way for users to attach meanings to their tags. A MOAT meaning refers to a Web resource and is part of the tagging; for instance, users could attach a meaning from DBPedia (<http://dbpedia.org/>) or any other resource they choose. Another approach to add meanings to tags, Extreme Tagging, is presented in [17]. Since a tag can have different meanings in different contexts, tagging tags is used in that approach to disambiguate the respective contexts. The underlying meaning of a tag can be revealed by means of another tag. Kim *et al.* [8] propose a system named SCOT, a semantic cloud of tags to represent the structure and semantics of tagging data and to allow the import and export among different platforms. Tagging activities are represented as a tag cloud, which includes user-tag and tag-tag relations. SCOT uses SIOC to describe site information, FOAF to represent agents (both humans and machines), and SKOS to represent tags and allow semantic relations among them.

Oren *et al.* [18] explore the meaning of semantic annotations but do not propose an explicit model. According to them, tagging expresses an unspecified relation between the resource and the tag. Thus, according to their position, making a complex statement about the real world is not possible, but only assigning tags, because of the lack of context.

5.2 Findings and Results

Our proposal builds mainly upon the work of Gruber [5] and Passant & Laublet [9]. While our approach is different from all other ones we are aware of, it remains widely compatible with existing algorithms and tools. Currently, we do not yet explore the

problem of import and export of tagging data across platforms as done by Kim et al. [8]. Our proposal allows tagging tags same as in the work by Tanasescu & Streibel [17], and defining hierarchical relationships as e.g. offered by Bibsonomy (<http://www.bibsonomy.org/>). In contrast to existing approaches, our model expands the scope of taggable objects to a much broader set than described in any previous work, and uses this expansion to support the collective construction of conceptual graphs involving relationships between tags. Furthermore, by means of these networks, we expect to facilitate the disambiguation of tags in a similar way than Yeung, Gibbins & Shadbolt [16] propose.

Through the experiments described in the evaluation section, we gained preliminary evidence that our approach can be used with minimal instruction by average users familiar with traditional tagging systems. Also, we can see that a relatively simple expansion of the current tagging meta-model facilitates (i) the construction of conceptual graphs and (ii) the inference of a hierarchy of tags and other meaningful relationships such as synonyms and antonyms based on standard mining techniques, which will be immediately useful for query expansion and disambiguation.

6 Conclusions and Future Work

We have proposed a minimal yet fundamental expansion of the meta-model of tagging in order to empower the construction of richer conceptual structures while keeping the ease and popularity of free tagging. Our model adds a semantic level to free-tagging in order to improve search and retrieval by means of the addition of new elements in the tagging operation. Such can be used to (i) build complex conceptual graphs that represent the underlying relations between tags, (ii) improve those networks by social mechanisms for convergence, (iii) use those networks to disambiguate tag meanings and for query expansion, and (iv) reduce the gap for deriving formal structures from tagging data for other purposes.

Our model proved to be feasible, even though the first prototype needs to be improved and complemented in order to be able to collect more data about tagging activities in agents, tags, and relationships other than just pairs of tags. Additionally, more tests are required to determine more precisely how our model can be used (i) to improve search and retrieval in STS and social convergence mechanisms, and (ii) to derive formal structures from tagging.

According to the results of the evaluation, some improvements are needed and would be useful to derive formal structures and improve search and retrieval:

- Semi-automatic consolidation of data, *i.e.*, tags and relationships, by means of normalization and disambiguation techniques to reduce lexical variations and suggest hypernyms and synonyms.
- Contextualization of taggable objects to avoid misinterpretations such as taking “avianFlu” as an airline instead of a disease. It would be also useful to collect more meaningful relationships.
- Allowing the characterization of relationships (such as symmetry and transitivity) could be very useful. However, instructing average users to use this feature properly could be very difficult as well.

In a nutshell, we hope that our expanded model will help to improve the performance of tagging systems while keeping up with their popularity and ease of use on a Web scale.

References

1. VanDerWal, T. Folksonomy. 2007 [retrieved 2009-04-02]; Available from: <http://www.vanderwal.net/folksonomy.html>.
2. Passant, A., Using Ontologies to Strengthen Folksonomies and Enrich Information Retrieval in Weblogs: Theoretical background and corporate use-case, in: International Conference on Weblogs and Social Media. 2007: USA.
3. Specia, L. and E. Motta, Integrating Folksonomies with the Semantic Web, in: The Semantic Web: Research and Applications. 2007, Springer. pp. 624-639.
4. Hepp, M., E-Business Vocabularies as a Moving Target: Quantifying the Conceptual Dynamics in Domains, in: Knowledge Engineering: Practice and Patterns. 2008, Springer: Berlin / Heidelberg. pp. 388-403.
5. Gruber, T., Ontology of Folksonomy: A Mash-up of Apples and Oranges. International Journal on Semantic Web & Information Systems, 2007. 3(2).
6. Angelitou, S., M. Sabou, and E. Motta, Semantically Enriching Folksonomies with FLOR, in: Proceedings of the European Semantic Web Conference - Workshop on Collective Intelligence and the Semantic Web. 2008: Spain.
7. Golder, S.A. and B.A. Huberman, Usage patterns of collaborative tagging systems. Journal of Information Science, 2006. 32(2): pp. 198-208.
8. Kim, H.-L., et al., Social Semantic Cloud of Tag: Semantic Model for Social Tagging, in: Agent and Multi-Agent Systems: Technologies and Applications. 2008, Springer. pp. 83-92.
9. Passant, A. and P. Laublet, Meaning Of A Tag: A Collaborative Approach to Bridge the Gap Between Tagging and Linked Data, in: International World Wide Web Conference - Linked Data on the Web Workshop. 2008: China.
10. Gruber, T., Collective Knowledge Systems: Where the Social Web meets the Semantic Web. Web Semantics: Science, Services and Agents on the World Wide Web, 2007.
11. Lemieux, S. Social Tagging and the Enterprise: Does Tagging Work at Work? [retrieved 2009-02-10]; Available from: <http://www.semanticuniverse.com/articles-social-tagging-and-enterprise-does-tagging-work-work.html>.
12. Newman, R. Tag Ontology Design. 2004 [retrieved 2009-02-16]; Available from: <http://www.holygoat.co.uk/projects/tags/>.
13. Sabou, M., M. d'Aquin, and E. Motta, SCARLET: SemantiC relAtion discoverY by harvesting onLinE ontologies, in: European Semantic Web Conference. 2008: Spain.
14. Gracia, J., et al., Querying the Web: A Multi-ontology Disambiguation Method, in: International Conference on Web Engineering. 2006: USA.
15. Sabou, M., M. d'Aquin, and E. Motta, Using the Semantic Web as Background Knowledge for Ontology Mapping, in: International Semantic Web Conference - Workshop on Ontology Matching. 2006: Grecia.
16. Yeung, C.A., N. Gibbins, and N. Shadbolt, Understanding the Semantics of Ambiguous Tags in Folksonomies, in: International Workshop on Emergent Semantics and Ontology Evolution. 2007: Korea.
17. Tanasescu, V. and O. Streibel, Extreme Tagging: Emergent Semantics through the Tagging of Tags, in: International Workshop on Emergent Semantics and Ontology Evolution. 2007: Korea.
18. Oren, E., et al., What are Semantic Annotations. 2006.