Extreme Tagging: Emergent Semantics through the Tagging of Tags

Vlad Tanasescu[1], Olga Streibel[2]

[1] Knowledge Media Institute, The Open University,
Walton Hall, Milton Keynes, MK7 6AA, UK
v.tanasescu@open.ac.uk, vladtn@gmail.com

[2] Netzbasierte Informationssysteme, Freie Universität Berlin,
14195 Berlin, Germany
streibel@inf.fu-berlin.de, ostreibel@gmail.com

Abstract. While the Semantic Web requires a large amount of structured knowledge (triples) to allow machine reasoning, the acquisition of this knowledge still represents an open issue. Indeed, expressing expert knowledge in a given formalism is a tedious process. Less structured annotations such as tagging have, however, proved immensely popular, whilst existing unstructured or semi-structured collaborative knowledge bases such as Wikipedia have proven to be useful and scalable. Both processes are often regulated through social mechanisms such as wiki-like operations, recommendations, ratings, and collaborative games. To promote collaborative tagging as a means to acquire unstructured as well as structured knowledge we introduce the notion of *Extreme Tagging*, which describes systems which allow the tagging of resources, as well as of tags themselves and their relations. We provide a formal description of extreme tagging followed by examples and highlight the necessity of regulatory processes which can be applied to it. We also present a prototype implementation.

Keywords: semantic web, web2.0, tagging, emergent semantics, meaning, semantic associations, knowledge paths.

1 Introduction

The process of building “a new brain for humankind” [1] as foreseen by semantic web research appears to be a slow one. Indeed, the semantic web contributed to the success of the notion of *ontology*, “a logical theory accounting for the intended meaning of a formal vocabulary, i.e. its ontological commitment to a particular conceptualisation of the world” [2], but, possibly due to lack of software support [3], ontologies are difficult to build, even at the community level. Moreover, the final aim of the semantic web – data integration through ontology matching – is still a research question as it can be automated only in simple cases. Indeed, although there are already a large number of RDF files on the web, whether manually or automatically generated, only about 25 000 documents [4] representing semantic models, i.e. ontologies, are available online. This should not be the case, as ontologies should be easy to produce by each community, then shared in order to be aligned with others using the stack of specifications and languages – the semantic web "layer cake" [5] – designed to support this task [6].

In practice however, building and matching ontologies, appears to be an expert task, and difficulties related to knowledge acquisition, experienced decades ago in the artificial intelligence community, resurface. Moreover, while ontologies seem well suited to the description of scientific domains such as medicine and biology which are already semi-formal and organized by categories and part-of relationships, some communities such as geospatial scientists only accept with scepticism the exclusive usage of ontologies to describe their domains [7]. Arguments in favour of using alternative knowledge representation models include, amongst other, the inadequacy of category based reasoning to represent reality [8], the absence of grounding of symbolic systems [9], the need of different representations of the same entity according to the context [10], as well as the difficulty to represent psychological concepts such as *affordances* in a hierarchical way [11]. Indeed, we are still waiting for ontologies to be flexible enough to match the representational complexity of the human mind.

In the meantime, so called Web2.0 applications, by motivating users to contribute information, introducing fine tuned social regulation mechanism, as well as providing friendly user interfaces, have been experiencing both phenomenal growth and success. With the advent of Web2.0 the usage of unstructured annotations such as *tagging*, spread widely. Although the relation of tagging and social interaction has not, to our knowledge, been investigated in the literature, it seems to be the only way to allow users to describe their own content, since the system cannot determine in advance what this content will be. Collaborative tagging systems, by renouncing the use of predefined vocabularies, provide a simple way for users to give their own meaning to their own content [12].

Therefore, while current research is still trying to alleviate problems related to the practical use of ontologies, the semantic web may benefit from techniques used by Web2.0 applications. We believe that for the semantic web to expand faster, new semantic acquisition approaches, distinct from the centralized ontology development by experts, need to be explored. We also believe that any successful solution will use the social lever which raised the Web and Web2.0 to that level of popularity and usage.

Therefore, we introduce the notion of *Extreme Tagging Systems* (ETS), as an extension of collaborative tagging systems allowing the collaborative construction of knowledge bases. An ETS offers a superset of the possibilities of collaborative tagging systems in that they allow to collaboratively tag the tags themselves, as well as relations between tags. Unlike previous research on emergent semantics of collaborative tagging systems, ETS are not destined to exclusively produce hierarchical ontologies but strive to allow the expression and retrieval of multiple nuances of meaning, or semantic associations. The production of relevant semantic associations can then be automatically controlled through social network regulation mechanisms.

We first describe collaborative tagging systems. Then show the modifications introduced by extreme tagging systems, providing a formal definition. Accordingly, we explain our prototype implementation, and, before concluding, give some examples of regulation mechanisms that should be applied to the system.

2 The Semantics of Collaborative Tagging Systems

Collaborative tagging systems (CTS) support multiple users in the activity of tagging, which is marking content for future navigation, filtering or search [13]. As there is no prior agreed structure or shared vocabulary CTS users need neither prior knowledge nor specific skills to use the system [14]. We prefer to avoid the term folksonomy [15], not only because it is ambiguous (as stated in [13]), but also because of the relation to taxonomy, which seems to us unjustified in that context.

Tagging systems can be represented as hypergraphs [16] where the set of vertices is partitioned into sets:

$$U = \{u_1, \ldots, u_k\}, R = \{r_1, \ldots, r_m\}, \text{ and } T = \{t_1, \ldots, t_l\} \ . \quad (1)$$

U, *R*, and *T* correspond to *users*, *resources*, and *tags*. An annotation, i.e. a resource tagged with a tag by a user, is an element of set A, where:

$$A \subseteq U \times R \times T \ . \quad (2)$$

The final hypergraph formed by a collaborative tagging system is defined as *G* with:

$$G = \langle V, E \rangle \text{ with vertices } V = U \cup R \cup T, \text{ and edges } E = \{\{u, r, t\} \mid (u, r, t) \in A\} \ . \quad (3)$$

Collaborative tagging systems have proved extremely popular. Their strengths consist in generating *serendipity* while browsing – the fact of being able to retrieve what others have tagged in a similar way, e.g. one can retrieve everything that has been annotated using the tag “ant” –, as well as the elaboration of *desire lines* – a non constrained reflection of the user’s vocabulary – through a dataset [12] (e.g. I can use the English tag “ant” or the French “fourmi” indifferently, without being constrained by the system). However, when compared to more formal descriptions of domains, CTS are criticized for their *ambiguity* (an “ant” tag may be found for a resource related to “Actor Network Theory”, the “Apache Ant project”, or a representation of the insect), the dealing with *multiple words* constituting a single tag (“semantic web”, “semanticweb” or “semantic-web” for example) or *synonymy* (“mac” “macintosh”, and “apple”). These issues have leaded some to colloquially describe tagging systems as “a mess”.

To go toward “less mess”, approaches have been proposed to find groups of related tags by using tag co-occurrence for given resources [17][18][19]. Moreover, most websites using collaborative tagging systems already present tag clouds – a representation of a resource’s annotations where each tag is visually weighted by his number of occurrences –, or allow presentation by tag clusters – several tags are grouped under an appellation – and often offer tag recommendations – tags are suggested according to previous annotations.

Furthermore, some semantic web oriented approaches attempt to extract ontologies from collaborative tagging systems. In [16] the author maps tags onto concepts and resources to instances and applies network analysis techniques to cluster them. [20] presents an ontology for tags which would allow them to be shared and exchanged between systems, while in [21] the authors mine association rules between tagged resources to recommend tags, users, or resources, discovering supertag relations as well as resource communities. In [22] the authors deduce clusters and relations between tags by relating them to background knowledge obtained through ontology searches while [14] presents an experiment to automate the previous method.

Ultimately, ontologies and tagging systems are both symbolic frameworks, and as such they are subject to the criticism of the lack of a retrievable grounding. Indeed, in both cases by using symbols (in a given language), the expressed concept, or *signified*, remains in people's minds, and the resulting symbol networks may appear – especially to a machine – as "free floating island[s] of reeds [with] no anchor in reality" [9]. However, CTS usually tag existing resources, i.e. specifying the referent, or ground, that symbols denote, without indicating the details of this denotation, as opposed to ontologies which first have to describe a domain before adding instances, and limit the grounding to a few pre-existing relations, i.e. the ones defined in the ontology (e.g. "part-of") plus the one assumed by the model (e.g. "is-a", "subclass-of", etc). Extreme tagging, by allowing the tagging of tags as resources as well as the specification of the relations between tags, is an attempt to push symbolic annotation frameworks to an extreme in order to see what the grounding problem becomes when any relation can be symbolically described at an arbitrary level of granularity.

3 Extreme Tagging Systems

An Extreme Tagging System (or ETS) offers a superset of the possibilities of collaborative tagging systems in that it allows to collaboratively tag the tags themselves, as well as relations between them. For example, a media resource representing the close-up of a car may be tagged with "car", "wheels" and "travel". The tag "wheels" itself may then be tagged (possibly by a different user) with "car" and "wheel", and the tag "car" itself could further be tagged with "vehicle" (cf. Figure 1)[1].

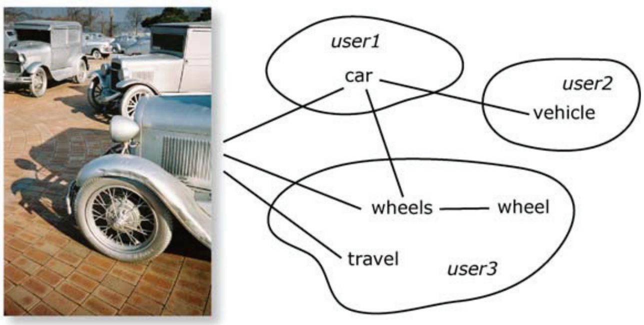

Figure 1. Tagging the tags.

[1] Picture from Flickr user Anjuli: http://www.flickr.com/photos/49502989227@N01/56641591/

The tagging of tags is justified by the fact that a tag can have different meanings in different contexts: tagging the tags and the relations between them is used to disambiguate these contexts. For example the tag "tank" on the Flickr photo sharing service[2] is used to tag military vehicles[3], fish tanks[4] as well as a person[5]. Tag tagging allows a user to explain the meaning of her or his annotations. It also reveals the multiciplicity of meanings: by tagging "tank" with "fish" *and* "vehicle", the ambiguity becomes apparent and users can then decide to filter accordingly.

The operation of tagging introduces a relation which is not only functional (something has been tagged by somebody) but also a *meaningful* (for some reason, excluding spam, somebody tagged something with this particular tag). Indeed tagging a picture with "wheels" may relate to what the picture depicts, and a tag "travel" may relate to the origin of the picture. However, the meaning of the relation is not made explicit by the user at the moment of tagging: we believe that not having to think precisely to the relation and verbalise it as one would do in an ontology results in a smaller cognitive load for the user and is part of the appeal of tagging. In extreme tagging however, this relation itself can be tagged, later on, by any user. The operation of tagging relations between tags can naturally be expressed by triples, for example, if "_" represents the *implicit* relation introduced by the tagging operation itself, while "…" is used to represent *any tag*, relations can be: (resource, {_}, "wheels"), (resource, {_}, "travel"), (resource, {_}, ...), (resource, {"shows"}, "wheels"), (resource, {"represents"}, "wheels"), (resource, {...}, "wheels") or (resource, {"taken-during"}, "travel"), etc. (cf. Figure 2)

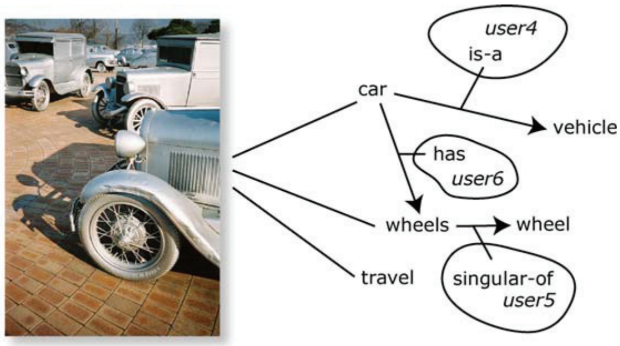

Figure 2. Tagging relations between tags.

Allowing users to tag the tags and the relations between them leads to the generation of *Semantic Associations*. Semantic Associations are chains of relations between one tag to another, or, in graph theoretic terms, a labelled path between two nodes. According to the definitions of [23] and [24], two entities are semantically associated

[2] Flickr, http://www.flickr.com/
[3] e.g. http://www.flickr.com/photos/barryslemmings/tags/tank/
[4] e.g. http://www.flickr.com/photos/towert7/tags/tank/
[5] http://www.flickr.com/photos/50836387@N00/tags/tank/

if they are *semantically connected*, i.e. there exist a path of relations between them, or *semantically similar*, i.e. two entities are similar if a path from the first one to another is similar to the path from the second one to another. We also call semantic annotations *knowledge paths*, as in this context they represent a crystallisation of the users' knowledge. We consider that the tagging relation itself, even if implicit, qualifies as a relation in a knowledge path, while we consider that the notion of semantic similarity can be extended from subclass/superclass relations only to any similarity measure. Collaboratively tagging resources, tags and relations leads to serendipitous discovery of associations between resources and/or tags. An example path between "wheel" and "vehicle" for example would be, expressed as a list of triples <"wheel", "vehicle"> = [("wheel", {"singular-of"}, "wheels"), ("wheels", {_}, "car"), ("car", {"is-a"}, "vehicle")].

The ETS model is defined as a collaborative tagging system with semantic associations. Therefore ETS are extensions of the formal model for collaborative tagging systems, defined as follows:

$$\Omega = \langle U, T, A, D \rangle \text{, where } A \subseteq U \times T \times T \text{ and } D \subseteq U \times T \times T \times T \quad . \tag{4}$$

We do not distinguish between the set of resources/entities *T* and the set of tags: all elements of *T* are entities, which can be "tags" or "resources". Indeed the mapping description of each entity by a unique identifier – in practice, a URI – makes the distinction superfluous. *A* is the set of assignments, as in traditional CTS while *D* represents directional annotations of relations between entities (tags or resources). According to this definition an ETS becomes a hypergraph:

$$G = \langle V, E \rangle \text{, with vertices } \mathrm{V} = \mathrm{U} \cup \mathrm{T} \text{, and edges} \tag{5}$$
$$E = \{\{u, r, t, d\} | (u, r, t) \in A \vee (u, r, t, d) \in D\} \text{ .}$$

The distinction between *A* and *D* reflects the distinction between *implicit* and *explicit* relations. An implicit relation occurs when an entity has been tagged while an explicit one appears if the relation between two entities has itself been tagged. A knowledge path is a path consisting of explicit or implicit relations between entities.

As relations between tags constitute triples, the link to RDF becomes obvious. Indeed ETS have the same goals as those sometime advocated by RDF proponents, "to allow anyone to say anything about anything" [25]. However, if ETS triples can be represented as RDF, extreme tagging introduces novelties. Indeed, RDF resources acquire their unique identity through the use of namespaces which contributes to slowing the process of knowledge acquisition as pre-existent knowledge about entities is needed. For example in the context of fish tanks the entity "http://fish.com/#tank" is needed, instead of "http://military.org/#tank". In ETS however, a tag is tagged by all its meanings, and disambiguation occurs during the query process, not at the tag description level, i.e. "tank" is only one tag, with a unique URI. If it is tagged as container and as a weapon, disambiguation will occur during knowledge path elicitation, as the knowledge path leading from "tank" to "fish" or "sea" will only use one of the meanings.

4 Tagopedia: an Extreme Tagging System

Tagopedia[6] is a prototype ETS built on top of the Facebook platform. Facebook is a social network web application providing a developer framework allowing the creation of applications which interact with core host features such as profile management and login. As any collaborative tagging system Tagopedia allows to tag resources, represented by URIs (cf. Figure 3).

Figure 3. Basic Tagopedia usage.

When clicking on a tag however, the user is asked to tag the chosen relation or to enforce an already existing relation by selecting it (cf. Figure 4). The application then moves to the target tag, showing the linked tags and resources and allowing to define new relations. The user may also choose not to tag the relation and directly reach the target, keeping it implicit.

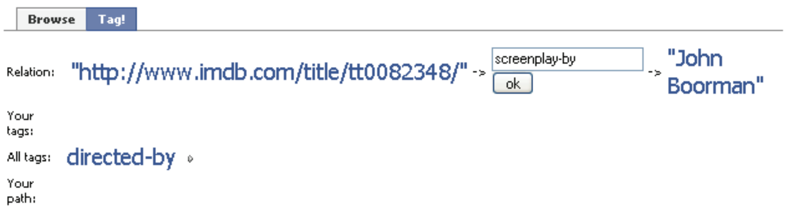

Figure 4. Tagging relations in Tagopedia.

Here is an example of collaboratively build semantic associations between entities in Tagopedia, written as a list of triples:

sa1:
[("John Boorman", {directed}, "Excalibur"),
("Excalibur", {about}, holy-grail) ,
(holy-grail, {similar-to}, grail),

[6] Available at http://apps.facebook.com/tagopedia/ (a Facebook account is required). The name *Tagopedia*, proposed independently by the authors, has already been proposed in 2005 by Russell Beattie in a blog post, for a related application (http://www.russellbeattie.com/notebook/1008277.html).

(grail, {topic-of}, "http://www.imdb.com/title/tt0071853/"),
("http://www.imdb.com/title/tt0071853/", {has-title}, "Monty Python and the Holy Grail"),
("Monty Python and the Holy Grail", {directed-by}, "Terry Gilliam")]

sa2:
[("John Boorman", {is-a}, film-director), (film-director, {includes}, "Terry Gilliam")]

sa3:
[("John Boorman", {is}, British), (British, {nationality-of}, "Terry Gilliam")]

5 Emergent Semantics

In ETS, semantics are related to the users' activity and input. The operations involved in a user's activity can be classified as: *annotation*, *navigation* and *control*. At each level there is a need for *incitation*, a means to motivate the user to use the system. As a result of these three operations, tags are created and annotated collaboratively and unconstrained semantics emerge (cf. Figure 5). In this section, we describe each activity in turn as well as the corresponding motivation mechanism:

Figure 5. Plurality of meanings.

Through *annotation* users are given the opportunity to create their personal knowledge base. Indeed, instead of tagging resources at their hosting websites, building unrelated islands of tags, they can relate all their resources with their own meaning. Moreover, they can access, for the same resource, tags from other users, and decide to explore their meaning by navigating to them.

Through *navigation*, users build or enforce semantic associations. Indeed, by exploring a tag which tags a tag, either the user is looking for an explanation of this tag, or she already knows the relation. We assume that she knows the relation if a) she tagged it before, or b) she chooses to tag it when asked to do so. As previously mentioned, navigation does not happen between a tag and another tag without presenting the relation, which the user can choose to tag or not. The motivation of this additional step is to constrain the meanings obtained. Paths which have been explored and validated, are recorded and displayed the next time a request is made to find the paths from one node to another.

Finally *control* mechanisms are necessary in order for the system to evolve, some of these control mechanisms can be:

1) *total control over ones annotations*: the annotations added by a user can be modified or deleted by her.
2) *appreciation and depreciation of tags*: a user can rank a tag (+ or – only). If the total ranking goes below a given threshold, the tag becomes "private" and does not appear in public searches any more. A similar method is already used by commercial websites[7].
3) *questions to author*: Facebook, just as other social networks provides the notion of "friends", or "contacts", i.e. users which acknowledged a mutual relationship. If a user does not understand a tagging made by one of her related users a quick means is provided to send him or her a message to ask for an explanation, i.e. the tagging of this particular relation. The requester is notified as soon as the explanation has been given.

It is assumed that each user is interested in sharing his or her vision of the world and in discovering other ways of perceiving it. To the first interest corresponds the *annotation* activity as well as some control activities number 1) and 2). As a further incitation annotating increases the user's ranking, in a similar way as internet forums display titles according to the number of posts (often quite imaginative, for example using a graduated scale going from *rookie*, to *half-god* or *absolute guru*). In parallel, an increase in status can be achieved through *navigation* only, in a similar manner to some multiplayer computer games which increase the avatar's status by providing titles according to the percentage of the virtual map explored. Indeed, These two ways of using the system, annotating and creating, combine when a *navigator* – i.e. a user who mostly navigates the system, comparable to a Wikipedia reader rather than to an editor – earns creation points by completing paths and *creators* earn more points if the paths they created are navigated (i.e. if they make sense). Further incitation may involve *visualisation* of the number of elements created, as well as graph presentation of the paths explored.

6 Conclusion and Future Work

The benefits of pushing tagging to an extreme are the ease with which knowledge is acquired, as well as the comprehensiveness of the resulting KB. Possible caveats, which we believe can be solved by collaborative means, include the difficulty to assess the relevance of the resulting knowledge in a given context. Tagopedia is a first prototype of an Extreme Tagging System and we are waiting to obtain a larger knowledge base to attempt a serious evaluation. However, we used the prototype in a limited environment composed of 5 users, and, from the amount of serendipitous meaning collected, were already convinced of the interest of the system. We are planning to release it to the Facebook community in the following months and explore the

[7] e.g. Spockcom, http://www.spock.com/.

aforementioned control mechanisms through it. We are also working on an RDF export mechanism as well as on the integration of a SPARQL query engine. We also plan to import large amounts of tags from Wikipedia and other websites, using links inside the pages or other structured information in order to populate the knowledge base.

Acknowledgments. The authors are grateful to the anonymous reviewers for their precious comments. They would also like to thank Yiwen Wang, Kaixuan Wang and Fadi Badra, who contributed ideas and energy in the early stages of this project. Further thanks are extended to Sean Bechhofer, Enrico Motta, and John Domingue, who contributed, during SSSW07 and later on, to the success of this enterprise, and to Lyndon Nixon for proofreading this paper.

References

1. D. Fensel and M. Musen, "The semantic web: a brain for humankind", Intelligent Systems, vol. 16, pp. 24-25, 2001.
2. N. Guarino, Formal Ontology in Information Systems: Proceedings of the 1st International Conference June 6-8, 1998, Trento, Italy. IOS Press, 1998.
3. M. Dzbor, E. Motta, C. Buil, J. Gomez, O. Görlitz & H. Lewen: Developing ontologies in OWL: An observational study. OWL: Experiences & Directions Workshop, Georgia, US, 10-11 November 2006
4. M. d'Aquin, M. Sabou, M. Dzbor, C. Baldassarre, L. Gridinoc, S. Angeletou, and E. Motta. WATSON: A Gateway for the Semantic Web. Poster session of the European Semantic Web Conference, ESWC 2007.
5. The Semantic Web "Layer Cake": http://www.w3.org/2004/Talks/0412-RDF-functions/slide4-0.html
6. Sir Tim Berners-Lee at the Oxford Internet Institute, from webcast: http://webcast.oii.ox.ac.uk/?view=Webcast&ID=20060314_139
7. W. Kuhn: Why Information Science needs Cognitive Semantics - and what it has to offer in return, 2003.
8. A. Wierzbicka: "Apples" Are Not a "Kind of Fruit": The Semantics of Human Categorization American Ethnologist, JSTOR, 11, 313—328, 1984.
9. P. Gärdenfors: How to Make the Semantic Web More Semantic Formal Ontology in Information Systems, Proceedings of the Third International Conference (FOIS), 17-34, 2004.
10. C. Vangenot, C. Parent and S. Spaccapietra: Modeling and manipulating multiple representations of spatial data Proc. of the Symposium on Geospatial Theory, Processing and Applications, 2002.
11. J. Gibson: The Ecological Approach to Visual Perception, Lawrence Erlbaum Associates, 1979, p.42-43
12. A. Mathes: Folksonomies-Cooperative Classification and Communication Through Shared Metadata Computer Mediated Communication, LIS590CMC (Doctoral Seminar), Graduate School of Library and Information Science, University of Illinois Urbana-Champaign, December, 2004
13. S. Golder & B. Huberman: The Structure of Collaborative Tagging Systems, Arxiv preprint cs.DL/0508082, 2005
14. Angeletou, S., Sabou, M., Specia, L., Motta, E., (2007) Bridging the Gap Between Folksonomies and the Semantic Web: An Experience Report. Workshop: Bridging the Gap between Semantic Web and Web 2.0, European Semantic Web Conference.

15. Vander Wal, T. Folksonomy, http://www.vanderwal.net/folksonomy.html, 2007
16. Mika, P. Ontologies are us: A unified model of social networks and semantics Proc. ISWC2005, 2005
17. P. Schmitz. Inducing Ontology from Flickr Tags. In Proc. of the Collaborative Web Tagging Workshop at WWW'06, 2006.
18. X. Wu, L. Zhang, and Y. Yu. Exploring Social Annotations for the Semantic Web. In Proc. of WWW'06, 2006.
19. G. Begelman, P. Keller, and F. Smadja. Automated Tag Clustering: Improving search and exploration in the tag space. In Proc. of the Collaborative Web Tagging Workshop at WWW'06, 2006.
20. Gruber, T., Ontology of Folksonomy: A Mash-up of Apples and Oranges. AIS SIGSEMIS Bulletin, 2005. 2 (3&4).
21. Schmitz, C.; Hotho, A.; Jaschke, R. & Stumme, G. Mining association rules in folksonomies Data Science and Classification: Proc. of the 10th IFCS Conf., Ljubljana, Slovenia, July, Springer, 2006
22. L. Specia and E. Motta. Integrating Folksonomies with the Semantic Web. In Proc. of ESWC'07, 2007.
23. A. Sheth et al., Semantic Association Identification and Knowledge Discovery for National Security Applications. 2004.
24. K. Anyanwu and A. Sheth, "P-Queries: enabling querying for semantic associations on the semantic web", Proceedings of the 12th international conference on World Wide Web, pp. 690-699, 2003.
25. I. Davis, Introduction to RDF slides, http://research.talis.com/2005/rdf-intro/#(7), 2005