The Automatic Integration of Folksonomies with Taxonomies Using Non-Axiomatic Logic

Joe Geldart and Stephen Cummins

Durham University

Cooperative tagging systems such as folksonomies are powerful tools when used to annotate information resources. The inherent power of folksonomies is in their ability to allow casual users to easily contribute ad-hoc, yet meaningful, resource metadata without any specialist training. Folksonomies are becoming a popular 'Web 2.0' technology and are becoming more and more prevalent in education and knowledge management domains. Despite the wide ranging success of folksonomies they do exhibit some problems such as the inability to represent structured knowledge or semantics. Folksonomies fail to capture relationships between tags and indeed any semantic meaning signified by tags. Furthermore, older folksonomies have begun to degrade due to the lack of internal structure and from the use of many low quality tags. This paper describes a remedy for some of the problems associated with folksonomies. We introduce a method of automatic integration and inference of the relationships between tags and resources in a folksonomy using Non-Axiomatic Logic. We test this method on the CiteULike corpus of tags by comparing precision and recall between it and standard keyword search. Our results show that Non-Axiomatic reasoning is a promising technique for integrating tagging systems with more structured knowledge representations.

Introduction

This paper addresses some limitations inherent in the usage of community-generated tags (or keywords) for electronic resource annotation. Cooperative tagging systems allow their users to annotate resources with short, pithy tags which are relevant to themselves without enforcing a universal vocabulary. This allows users to build up a personal system that facilitates information retrieval as well as allowing searches across all users' tags. The word folksonomy is a portmanteau of the word folk, meaning people, and taxonomy, meaning a system of classification. The word was originally coined by the Thomas Vander Wal in August 2004. Folksonomies are a particular type of community tagging system and are essentially comprised of a flat classification system constructed out of the raw tagging data. Each resource is simply classified as being within the classes to which each tag implicitly denotes.

There has been a large increase in the use of folksonomies in recent years; in part this is due to the popularization 'Web 2.0' technologies, of which folksonomies are an example. Folksonomies are often considered as a 'grass-roots' approach to semantic indexing; decentralizing and democratizing information system organization.

The lack of central control, while an asset in some situations, can lead to problems when the system's users disagree about the tagging of resources. Existing tagging systems treat tags as an all-or-nothing affair with folksonomies based on classical classification. We argue in this paper that this is partially responsible for the marked deterioration of the quality of large-scale tagging systems.

Efforts have been made to integrate formal, 'expert' knowledge representations with community-generated folksonomies as part of a solution to the problem of quality. However, current approaches still suffer from the problem of disagreement.

This paper proposes a solution to the problem of integrating ontological information with folksonomies through the use of a logic with an evidence-based semantics. The main contributions we make are:

- An overview of the problem of integration and folksonomy quality (section 2).
- A solution to this problem using Non-Axiomatic Logic (section 3).
- Experimental evidence showing the reasoning capabilities of Non-Axiomatic Logic with tagging system data (sections 4 and 5).

We conclude with a survey of related research and a discussion of future research directions for non-axiomatic tagging systems.

The Problem of Integration

A problem for traditional tagging systems is that individual tags do not have explicitly defined relationships with each other. That is, there is no explicit classification mechanism between tags within a collaborative tagging system or folksonomy. This leads to further limitations, such as natural language features not being considered in resource searches. For example, the use of plurals, synonyms and homonyms may result in some resources not being retrieved when using standard lexical searches. The lack of semantics captured by simple metadata such as tags leads to an inability for artificial agents to perform advanced searching through use of logical inferences.

The limited ability of computational agents in searching these systems derives from the lack of semantic metadata implicitly defined in folksonomies. This is considered as being one of the main weaknesses of collaborative tagging technologies (Specia & Motta, 2007). Providing agents with the ability to perform logical inferences within collaborative tagging systems would bring such systems in line with other more formal mechanisms of resource management, such as ontologies. The major benefit of these tagging systems, which inevitably is the reason why they have been so successful, is the inherent flexibility and the overall ease of the tagging process. If this can be combined with a means of automatically capturing the semantic and logical relationships between tags whilst preserving the flexibility and convenience, it will be a much more versatile tool.

Recent literature presents solutions that apply manual or semi-automatic methods for integrating semantic metadata with folksonomies. CommonFolks is one such system that is essentially a hybrid of a folksonomy and ontology. This system does allow the richer more structured metadata commonly associated with ontologies to be associated with resources stored in a folksonomy. However, it shares the significant drawbacks of all such systems. The primary problem with existing hybrid approaches is that there is a change from being a flexible, quick and easy to use system to being a rigid system, with strict rules and requirements. These systems often require more time and effort from its everyday users in order to be used properly. It is arguably the freedom of expression and ease of tagging resources that makes folksonomies so successful, any hybrid system that directly or indirectly removes these benefits will no doubt result in a less popular system. The ideal would be to enable a richer metadata for resources to be generated purely from whatever information (tags) the users care to give. Thereby making use of what is already being put into the system instead of requiring more work from its users. This paper will go on to outline a proposed solution that aims to achieve this ideal.

The benefits of folksonomies are still being researched, many areas such as e-commerce [Amazon™], education [citeulike.org] and social networking [facebook.com] have displayed characteristics inherent in folksonomies. This indicates that folksonomies are still popular techniques for metadata generation and are becoming ever more widespread. Therefore addressing problems and improving performance is an important issue to investigate.

In an educational context specifically folksonomies fail to intelligently represent disagreement between users. In some educational and research disciplines there are various different and often conflicting perspectives on the same subject. These disagreements are often considered valid and indeed valuable to the discipline as a whole. In particular, sociology teaches that learning different theoretical perspectives is fundamental to the field. Even in the physical sciences there are some topics which hold a lack of consensus, such as the origins of the universe. The arguments for all theories should be represented in a knowledge management system. A knowledge system that does not cater specifically for the disagreement of its users is surely not a realistic or accurate representation of the knowledge domain. Since folksonomies are simple, flat classification systems they do provide a basic means for representing conflicting tags. However, since the tags have little semantics attached, the folksonomy itself cannot easily identify disagreements to its users; it is left to the users to manually find and identify resources that are tagged with conflicting tags. It is useful for the knowledge management system to identify topics that are contentious or under strong debate to its users, if only to educate the user to the fact that there is a disagreement in that particular topic.

It can be argued that there are very few subjects that have a simple right or wrong (true or false) answer. Allowing concepts to be expressed in different ways and encouraging debates, is considered as being beneficial to one's education. From a knowledge representation perspective there may be cases where disagreement is not desired, but even in these cases we believe that they can be found and handled more efficiently if the knowledge management system is able to automatically detect them and flag them to the user community for discussion. Either way, representing disagreement in large repositories of knowledge, like folksonomies, is an important consideration.

An Approach to Tagging with Non-Axiomatic Logic

As the previous discussion indicated, the main problems with the integration of ontological information with free-form folksonomic information are the flat semantics and potential for conflict of the latter. Previous attempts to integrate ontologies with folksonomies have focused on imposing a clear, universal semantics upon the latter; relegating disagreement to an aberration. As mentioned, this poses unique problems in educational and academic contexts where disagreement must be kept in order to faithfully represent the domain. This paper, therefore, takes the opposite position; opening ontologies to disagreement using a refined, evidence-based semantics.

A tagging system may be viewed as a collection of asserted pairs, each consisting of a resource and a tag. Each such assertion should have, from a logical perspective, a truth-value which may be used by reasoning mechanisms. Ideally, we want these truth-values to allow a spectrum of belief between true and false and the ability for inference rules to weight assertions by the amount of evidence that backs them. In other words, what we would like is for these truth-values to represent two different things; the degree to which the resource is judged to carry the tag, and the amount of user data which has gone into forming this judgment. A semantics with truth-values defined like this is known as an *evidence-based* semantics. An evidence-based semantics is explicitly partial, making explicit the fact that a user's assignment of truth-values is based upon a limited experience of the world.

In a Boolean truth-functional semantics, such as that used by first-order predicate logic or OWL, truth is taken to be a binary notion. A statement is taken to be either true or false, with no in between state. If a conflict arises between the truth-valuation of the statements of two users, there can be no middle ground and so one or the other of the statements must be thrown away.

Non-Axiomatic Logic (NAL) is a logic with an evidence-based semantics (Wang, 2005) that operates under an assumption of insufficient knowledge and insufficient resources (AIKIR) (Wang, 2006a); something the authors consider a good fit for the situation of the integration of ontologies and folksonomies. NAL is a term logic, as opposed to a predicate logic, with statements consisting of two terms (a subject and a predicate) linked by one of a small number of statement-forming primitive relations. The most important of these relations is the inheritance relation, which represents a (partial) sub-class relationship between the subject and the predicate.

```
S --> P
```
(1)

In (1), the term S is the subject whilst P is the predicate. The arrow represents the inheritance relation. The term S is said to be in the *extension* of P whereas P is in the *intension* of S. One of the features of NAL is that both extensional and intensional semantics may be represented in a common formalism.

Terms may themselves be built out of smaller parts. NAL defines a number of term-forming operators, generally available in both extensional and intensional forms (which cause different inferences to be made if the term is found in the subject or predicate position of a statement.) These operators include intersection, difference and a product-forming operator that constructs a tuple of other terms. This last operator addresses a common criticism of term logics; their inability to represent arbitrary relations between terms. There are also extensional and intensional image operators that allow statements to be made regarding the components of a product.

NAL also allows statements to be used as terms which, together with independent and dependent variables and implication and equivalence primitive relations, allow higher-order knowledge to be encoded; statements about statements. This permits the encoding of arbitrary structural inferences regarding a NAL knowledge base and even the representation of other logics.

The inference rules of NAL are based upon the notion of a syllogism (Wang, 2000). There are four syllogistic forms; deduction, induction, abduction and exemplification. The first is familiar from most introductions to logic, but the last three are maybe less so. Induction is the process of generalization from particulars, while its dual, abduction, produces particulars from generals. Exemplification is the dual to deduction and finds more specific terms from general ones. The combination of these four forms of reasoning in one logical formalism provides a powerful tool. From particular tagging patterns, NAL may be used to find relationships among tags by induction, applying such relationships via abduction. Particular hierarchies of tags can be used through deduction and exemplification. With the compound term constructors of NAL, complex classes (such as the rough equivalents on NAL terms to intersection, union and difference from set theory and many knowledge representation formalisms) may be automatically constructed. With the addition of products and higher-order sentences, arbitrary relational knowledge may be represented in the same formalism allowing, for example, the ability to develop refined knowledge based on resource metadata. Such a combination of these forms of reasoning is only possible in a semantics which allows assertions to be revised in a balanced fashion as they are rarely completely accurate.

Due to NAL's AIKIR assumption, a conventional reasoning algorithm would be inappropriate. NAL therefore uses a form of anytime reasoner, rather than a proof-theoretic or tableaux one, which is based upon this assumption. The Non-Axiomatic Reasoning System (NARS) is explicitly finite, and partial. Rather than attempting to reason over the entire knowledge-base simultaneously, it works on tasks (such as input judgments or queries) in a prioritized, probabilistic fashion. The finiteness of the reasoner's memory means that tasks or beliefs with low utility (those which are used very infrequently, or which haven't produced useful results) are forgotten when the system reaches its capacity limits. In this way, NARS not only avoids the issues of knowledge explosion which would otherwise result from having such a powerful and general inference process, but it also should cope well with poor tagging. Only resources and tags which are well-linked by inferencing and frequently referred to will persist. The system's judgment of 'well-linked' will depend on the contents of its background knowledge, and the system should become more resilient to such poor tagging as it gathers and infers more background knowledge.

The above is, necessarily, only a very brief introduction to the ideas behind NAL and NARS. The interested reader is recommended to read (Wang, 2006b) for a thorough introduction to the logic, its reasoning algorithm and its philosophical inspirations and implications.

Experiment

In order to test the efficacy of NAL as a logical basis for tagging systems we evaluated it against a simple, keyword-based system as a control. The evaluation criteria used is a measurement of the precision and recall of our solution in comparison with existing systems. This allows objective analysis of whether or not the solution presented in this paper improved the quality of searching in folksonomy-based systems.

The reasoner for the NAL-based tagging system is the open-source OpenNARS reasoner developed originally by Pei Wang and later converted by Joe Geldart into a reusable library. The control was written by Stephen Cummins and uses simple keyword indexing of resources.

The sample data for our experiment comes from CiteULike, a Web-based bibliography manager. CiteULike provides a complete corpus of anonymised tagging data extracted from its database in the form of a comma-separated variable file. The file contains a row for each tag attached to each resource, identified by article ID. Each row contains a field for the article ID and tag, together with a timestamp and anonymised user identifier. We prepared this data by reducing it to just the article ID and tag, and filtering out stop-words, commonly used words which are statistically insignificant on their own for purposes of semantic content analysis. The list of the stop words we used is included in appendix (Cummins & Geldart, 2008). We further reduced the size of the corpus to one hundred articles, randomly chosen but kept constant throughout the experiments.

The metrics chosen as useful indicators of system performance are precision and recall. Precision is defined as the fraction of the search result that is considered as being relevant for a particular query. Recall is considered as the ability of a particular search system to retrieve all relevant resources in its search results. This metric is typically difficult to measure in large search systems due to the fact that one requires knowledge of all the results retrieved as well as all of the results not retrieved. However, for the purposes of this paper we shall only be working with relatively small sets of data in order to demonstrate the concept.

We tested the two systems by comparing their precision and recall under a range of conditions. The precision and recall were calculated by dividing the tagging data in the corpus into two disjoint sets; input and test. The former set was used as input to the tagging systems, providing data about the resources. Queries consisting of random sets of tags were composed using both sets, and the precision and recall calculated. A system with high values for these metrics under this test would be able to predict the human-generated tags with high accuracy.

The experimental conditions consisted of varying the number of reasoning steps taken before querying (essentially, the length of the settling phase used to integrate the data with the reasoner) and the proportional size of the set of input data compared to the set of all data. It was expected that the OpenNARS-based reasoner would show positive correlation of both precision and recall with these conditions.

As NAL has an essentially different notion of membership of a result to a query result set than that of the classical-classification control, a notion of fuzzy cardinality was used for the precision and recall calculations. The two-component truth value obtained from OpenNARS was reduced to a one component expectation value, as described in (Wang, 1995). The sum of these expectation values across the result set was used to define the fuzzy cardinality. For the control, the expectation value for each result was one, capturing the classical notion of absolute classification. Such a calculation of precision and recall does slightly favor the classical system as each result from the NAL system counts for less (an expectation of one being impossible in NAL) and the results should be read with that in mind.

Evaluation

The experiment was run over a range of eleven settling lengths (from 0 to 100 in increments of ten steps, then from 100 to 900 inclusive with increments of 100) and ten different data proportion sizes. The results were then processed to produce statistics on the recall and precision rates of the two tagging systems.

The statistics are shown in Table 1. As can be seen, the NARS-based system shows (as compared to the control) a significantly higher mean recall rate but a significantly lower mean precision rate (across all conditions.) This, we believe, shows both the promise and the problems associated with the use of NARS in a tagging environment.

There were two main contributing factors to the low precision of the NARS-based tagging system. Firstly, the current version of OpenNARS (as of April

	Control		OpenNARS	
Precision	Mean	0.429	Mean	0.017
	Variance	0.245	Variance	0.002
Recall	Mean	0.393	Mean	0.681
	Variance	0.222	Variance	0.198

Table 1. Precision and recall statistics.

2008) was originally designed to be used as a stand-alone Java applet, as opposed to as a reasoner library. This has made the creation of an interface between the experimental harness and the reasoner fraught with difficulties, particularly with the appropriate decoding of statements from the reasoner. The authors decided to err towards over-production so that they might manually inspect interesting inferences from the system. The authors, having learnt from this experience, are working on a modular rewrite of OpenNARS which will allow for its use in other information systems. Secondly, the data-set we used consisted simply of a list of resource and tag pairs which we encoded as inheritance statements with frequency of 1. A more realistic system would likely include background knowledge (either through explicit inclusion of ontological information or simply by long-term inferencing on a larger set of taggings) which would reduce the judged truth of more spurious inferred taggings. A third possible, and we feel more interesting, reason for the low precision is that OpenNARS is generating entirely novel taggings which haven't been expressed by any user of CiteULike. A check of some of the inferences produced has shown a large number of correct (to the authors) taggings which aren't found in the corpus. This raises a question of the development of appropriate testing techniques for uncertain, intelligent systems, which the authors leave open for the moment.

The data was also analyzed to find the correlation co-efficients of the precision and recall to the experimental conditions. This is shown in Table 2. As can be seen, there is a very weak (anti-)correlation between the settling period and both precision and recall, for both reasoners. The very weak anti-correlation between precision and settling time for the NARS-based reasoner could be taken as a first sign of novel generation. This also points to a strategy for detecting novel generation; calculate their (inverse) correlation with the length of the settling period. As may also be seen, recall for NARS is most affected by the data size and to a degree greater than the control. We interpret this as an indication of the inherent growth of inferences as compared to input knowledge.

Control	Settling	Data size	Precision	Recall
Settling	1.0	6.551	0.0098	0.018
Data size	6.551	1.0	0.270	0.263
Precision	0.0098	0.270	1.0	0.961
Recall	0.018	0.263	0.961	1.0

OpenNARS	Settling	Data size	Precision	Recall
Settling	1.0	6.551	-0.011	0.059
Data size	6.551	1.0	-0.062	0.317
Precision	-0.011	0.270	1.0	0.193
Recall	0.059	0.317	0.193	1.0

Table 2. Correlation coefficients.

Related Work

Attempting to integrate the richness of semantic metadata with the ease of use exhibited by folksonomies is not a new concept as previously mentioned. There are studies that attempt to do this through using manual and semi-automatic techniques in order to enhance simple tagging based systems (Al-Khalifa & Davis, 2006; Bateman, Brooks, & McCalla, 2006; Bateman, Farzan, Brusilovsky, & McCalla, 2006). None of these systems however demonstrate a fully automated approach to generating semantic information simply from the tags used in the folksonomy. Similar research has been completed, however the focus is usually on modifying exiting folksonomies in order to operate in line with semantic web or ontology (Angeletou, Sabou, Specia, & Motta, 2007) based rules.

By far the most common method of integrating folksonomies with ontological metadata is to require users at the time of tagging to also manually position the resource in the ontology. This is the case in CommonFolks (Bateman, Brooks, & McCalla, 2006) where the users must define how the resource relates to other concepts in the WordNet™ ontology. Techniques exist that speed the whole process up for the users such as making automated ontology suggestions, however this still require a human to review them and make corrections, thus increasing the time overhead. This technique does improve the amount of semantic data recorded but reduces the ease of annotation and still does not specifically cater for disagreements in tags and indeed in relationships between other concepts in the ontology.

The resulting decrease in flexibility and speed of adding new resources to the system is certainly going to adversely affect likelihood of users participating. There are other similar systems outlined in the literature however CommonFolks is a particularly relevant example. Interested readers may wish to read (Auer & Dietzold, 2006; Bateman, Farzan, Brusilovsky, & McCalla, 2006; Laniado, Eynard, & Colombetti, 2007; Specia & Motta, 2007)

The older and more popular folksonomy systems have begun to experience problems caused by use of numerous low quality tags during resource annotation. Some of these problems become more apparent when looking at subjects that have commercial relevance where a form of tag spamming (Weinberger, 2005; Xu, Fu, Mao, & Su, 2006) is surfacing.

Folksonomies are designed to be somewhat resistant to current spamming techniques as they operate as virtual democracies, meaning each tag on a particular resource signifies a single vote. Multiple occurrences of a single tag in a particular resource count as multiple votes and therefore increase the recorded value of that particular tag. Usually knowledge representation systems based on folksonomies require a user to have an account in order to contribute to the tagging community. This is an attempt to safeguard against anonymous spammers trying to bias the popularity of particular resources or tags.

There is limited amount of literature that addresses the problem of intelligent searching within folksonomies. There is some solutions for the detection and use of natural language features such as homonyms (Bateman, Brooks, & McCalla, 2006). However only projects that integrate folksonomies with a natural language ontology such as the WordNet™ lattice have been able to address them.

At the time of writing this paper there has been no research into the use of non-axiomatic logics as method of representing and managing a folksonomy, especially a folksonomy that can specifically cater for disagreements between its users.

Conclusions

This paper has outlined an original method of generating ontological information simply by using the data already contained in a fully functional folksonomy. Our research indicates that NARS can be used to generate semantic relationships intelligently using nothing more than user-annotated tags as metadata. Therefore, no more information is required from the casual users of the folksonomy in order to use this system. NARS is also able to make inferences based on the information given by users and suggest meaningful tags for resources stored in the system.

The experiment outlined in this paper has shown that NARS can maintain an active knowledge representation of a folksonomy and make inferences based on its contents. In addition, our results have shown that by using NARS as the backend reasoner for folksonomy searches we yield an improved recall; albeit at the expense of precision. However, we feel that the low precision may be improved, as discussed previously.

With further investigation, the new metadata suggested by the reasoner could be evaluated by a human user in order determine whether the precision is as low as our experiments would seem to suggest. It is entirely possible that the reasoner has determined that a given resource should be annotated with a tag that a user has yet to devise. Subjective validation of tag quality from a group of users can determine whether or not the tag suggested by the reasoner is valid and of high quality. As already mentioned, the correlation of inferencing and settling time provides another possible test as to whether the changes of behavior of a system such as NARS are due to the processes of inference. As regards the problem of low quality tags degrading the performance of folksonomy searches, NARS is able to 'clean-up' or forget unpopular or less useful tags from the global search space. By having a separate repository used to store individual users' tags for particular 'favorite' resources, the personalization of folksonomy searching could be merged with a better, more intelligent global search mechanism in a fully automated way. This would maintain the personalization element of folksonomy systems and still improve the quality of the metadata shared between users.

The ability for NARS to represent disagreement suggests that the approach outlined in this paper would be useful to various disciplines where disagreement may be prominent or desirable. Such domains include education, academia, and the open Web, all of which lack a central authority to impose a single classification upon the world.

For the authors, the real potential with the use of NARS lies with its ability to deal with highly structured yet uncertain information such that disagreement and other forms of conflict and agreement can be handled subtly through the use of evidence measures. A more robust, rational Web would allow for more automation, and less manual intervention by users. As such, we feel that NAL is a good candidate representation language for the semantics of the Web.

References

Al-Khalifa, H. S., & Davis, H. C. (2006, Nov 17 - 21 2006). *FolksAnnotation: A Semantic Metadata Tool for Annotating Learning Resources Using Folksonomies and Domain Ontologies.* Paper presented at the The Second International IEEE Conference on Innovations in Information Technology, Dubai, UAE.

Angeletou, S., Sabou, M., Specia, L., & Motta, E. (2007). *Bridging the Gap Between Folksonomies and the Semantic Web: An Experience Report*. Paper presented at the 4th European Semantic Web Conference (ESWC 2007)

Auer, S., & Dietzold, S. (2006). *OntoWiki: A Tool for Social, Semantic Collaboration*. Paper presented at the 5th International Semantic Web Conference

Bateman, S., Brooks, C., & McCalla, G. (2006, 20/06/2006). *Collaborative Tagging Approaches for Ontological Metadata in Adaptive ELearning.* Paper presented at the Fourth International Workshop on Applications of Semantic Web Technologies for E-Learning, Dublin, Ireland.

Bateman, S., Farzan, R., Brusilovsky, P., & McCalla, G. (2006). *OATS: The Open Annotation and Tagging System*. Paper presented at the Third Annual International Scientific Conference of the Learning Object Repository Research Network

Cummins, S., & Geldart, J. (2008). Stop words file. Retrieved 29th April, 2008, from http://www.dur.ac.uk/s.a.cummins/Research/ISD2008/stop_words.txt

Laniado, D., Eynard, D., & Colombetti, M. (2007). *Using WordNet to turn a folksonomy into a hierarchy of concepts*. Paper presented at the Semantic Web Applications and Perspectives (SWAP)

Specia, L., & Motta, E. (2007). *Integrating Folksonomies with the Semantic Web*. Paper presented at the 4th European Semantic Web Conference 2007

Wang, P. (1995). *Non-Axiomatic Reasoning System: Exploring the essence of intelligence.* Indiana University.

Wang, P. (2000). Unified Inference in Extended Syllogism. *Applied Logic Series*, 117-129.

Wang, P. (2005). Experience-grounded semantics: a theory for intelligent systems. *Cognitive Systems Research, 6*(4), 282-302.

Wang, P. (2006a). The Logic of Intelligence. *Artificial General Intelligence*, 31-62.

Wang, P. (2006b). *Rigid Flexibility: The Logic of Intelligence* (Vol. 34): Springer.

Weinberger, D. (2005). *Tagging and Why it Matters*: Harvard University - Berkman Center for Internet & Society, Publication No. 2005-07.

Xu, Z., Fu, Y., Mao, J., & Su, D. (2006). *Towards the Semantic Web: Collaborative Tag Suggestions*. Paper presented at the 15th International World Wide Web Conference (WWW 2006)