Social Graphic Tagging for Semantic Metadata and a Case Study on Consensus Discovery

Jie Yang and **Mitsuru Ishizuka**
Graduate School of Information Science and Technology
University of Tokyo, Tokyo, Japan
yangj@mi.ci.i.u-tokyo.ac.jp, ishizuka@i.u-tokyo.ac.jp

Abstract

The lack of semantic metadata is becoming a barrier for the in-depth study and wide adoption of Semantic Web. At the same time, folksonomy draws more and more attention as a promising source of semantic metadata. By avoiding the use of "a-priori" agreements on ontology, which is the main feature for folksonomy, the interoperations among metadata from different users could be likely supported by dynamically constructed emergent semantics. In this paper the authors propose a concept model that supports metadata generation by extending the ideas from folksonomy. A semantic layer is specified in the model that comprises three types of semantics varying from simple to complex to support different kinds of semantic interoperations. An implementation of the model, Lego-Note, the folksonomy system featured by graphic tagging is introduced. The preliminary experiment is performed as a case study, and the qualitative analysis on social consensus discovery is exhibited.

1 Introduction

One key idea of the Semantic Web[Berners-Lee *et al.*, 2001; Fensel *et al.*, 2002] is to enrich the web with metadata[1] which describes resources such as web pages, documents, photos and real world objects in a machine understandable format. The common approach to the goal is: An ontology is defined somewhere first, then the web pages are annotated using the predefined ontologies manually ([Kahan and Koivunen, 2001; Handschuh *et al.*, 2001]) or (semi-) automatically ([Kiryakov *et al.*, 2003; Domingue and Dzbor, 2004; Kettler *et al.*, 2005]).

This approach, however, has its defects that prevent it from being widely adopted. Firstly, an ontology is usually defined by a group of experts in a centralized fashion. It has been noticed that the centralized ontology is too strict to be adopted by the majority of people in diverse situations. Secondly, in order to use an ontology, the contents of it must be understood by the user (e.g. the ontology engineer, the programmer or the ordinary web user). It is time consuming, especially when there are several different ontologies about the same domain. It turns out that the startup cost for using an ontology is too high, which might frustrate the general web user from exploring further. As a result, although quite a lot work has been done, there still is a lack of semantic metadata. As argued by [Huynh *et al.*, 2005], that turns out to be a barrier for in-depth study and wide adoption of Semantic Web technology.

The idea of emergent semantics [Karl *et al.*, 2004; Cudre-Mauroux *et al.*, 2006] provides a promising solution to the above problem. In [Karl *et al.*, 2004] the authors argue that the use of "a-priori" agreements on concepts, i.e. ontologies, is insufficient in ad-hoc and dynamic situations where the interacting parties do not anticipate all the interoperations. [Karl *et al.*, 2004] suggests that the users of a distributed system should be able to create and use their own metadata without any controlled vocabularies imposed by the system. To achieve the semantic interoperability among the inconsistent metadata from different users, negotiations can be conducted to reach agreements on common interpretations within the context of a given task. Originally the peer-to-peer system is the commonly referred to scenario for emergent semantics, whereas recently folksonomy is being considered as another valuable source.

The folksonomy or "folk taxonomy", is formed by a community, where each user tags the web resource using his own words and shares them socially. Folksnomy system (e.g. del.icio.us[2] and flickr[3]) succeed mainly in getting people involved with its simplicity and practicality. That is (1) no predefined vocabularies impose to the user. The user can tag the content with any arbitrary word. So that there is almost no startup cost needed, and (2) the folksonomy system provides immediate benefits to the user. For example, the user can find interested information by exploring the same tags he shared with the community or track the "hot" topics (e.g. the most tagged web pages) in a breeze. The more a user tags and shares, the better service returns. This is a benign circulation that the semantic annotation systems failed to support.

[1] Typically, metadata is defined as data about data, which might take various forms such as slot-value pairs, or category names in a classification. Whereas, when we mention "metadata", we actually mean the semantic metadata. The latter is a group of RDF statements, being the special kind of metadata. In this paper they are used interchangeably.

[2] `http://del.icio.us`

[3] `http://flickr.com`

As a kind of social metadata that organizes web information, folksonomy suffers the problem of ambiguity of semantics too, but as is discussed in [Michlmayr, 2005] folksonomy provides emergent semantics and the dynamic consensus can be achieved by negotiation to support semantic interoperability. Prior work includes [Mika, 2005] that tries to build ontologies for communities from folksonomy and [Wu *et al.*, 2005] where a probabilistic generative model is proposed to derive the emergent semantics of the (synonymous and ambiguous) tags. In addition, [Xu *et al.*, 2006] proposes a collaborative tag suggestion algorithm to identify most appropriate tags, while eliminating noise and spam.

In this paper, we propose a concept model for social metadata by extending the ideas of folksonomy then present our implementing system following the model. The preliminary experiment to achieve consensus (emergent semantics) is performed and the qualitative analysis is exhibited.

2 Tag Ontology: A Concept Model for Social Metadata

Metadata is defined as data about data. It's information classifing other specific information in more general terms. Traditionally, metadata is created by either the authors of the content or professionals who organize informational content as their job. By "social metadata" we emphasize that the metadata is generated and shared by normal web users in the open web without requiring a pre-defined formal ontology. The tag ontology is proposed to model the user's tagging behavior, organize the annotating result and to facilitate deriving emergent semantics from social medadata. The underlying consideration, detailed structure and features of the tag ontology are presented below.

2.1 Semantics of Tagging

To begin with, we discuss the tagging action and analyze its relation with the tag. We argue that the semantics embodied in the tagging action is much richer than that of the involved tags and that the tag just reveals some aspects of the tagging semantics. The former is the targeting semantics that should be captured and utilized by the system. For example, in the most common case the user by tagging a web page with "semantic web" tag means "the content of this page relates to semantic web". Whereas there are also cases where a user uses "todo" to denote the future work, or "sesame/SemanticWebApp" to represent that the web page is about project "sesame" and "sesame" is a Semantic Web application. From the examples, we can say that tag help the web user share his understanding of the web content explicitly in a machine processable way. Although the tagging semantics is still partially and vaguely conveyed.

By reviewing the recent work on folksonomy, a trend can be noticed that progressing along the semantic continuum [Uschold, 2003] ranging from implicit semantics that exists only in the heads of the people, to machine-processable formal semantics. To represent and utilize the implicit tagging semantics, at the very beginning only the semantics of the tag (keyword) is taken into consideration, as del.icio.us does. Then [Mika, 2005] proposed a tripartite model $A \times C \times I$ for folksonomy system, corresponding to the set of taggers (users who tag), the set of tags (words) and the set of the contents be tagged (e.g. the web pages). The tripartite model unveils the naturally existing linkages among the three parties and are used to build domain ontologies. More recently, [Gruber, 2005] suggested a tag ontology with $Tagging(Object, tag, tagger, source)$ as the core concept. It can be observed that more and more semantics embodied in the tagging action can be delivered as more and more powerful models proposed.

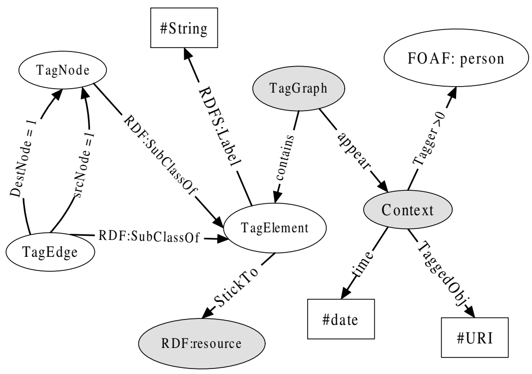

Figure 1: The structure of tag Ontology

The work in this paper follows the trend and proposes a tag ontology which is a step further along the semantic continuum towards the formal semantics.

2.2 Tag Ontology

Figure 1 presents the tag ontology. It is proposed to identify and formalize a conceptualization of the activity of tagging, and to support the communication among autonomous programs that commit it.

We model the tags attached to the web content as a graph (called tag graph) that takes the form of Direct Labelled Graph (DLG). The graph node is made up of keywords. Two nodes may get connected by the graph edge, to which arbitrary label can be attached. The action of attaching a tag graph to the web resource is called "graphic tagging". The left part of Fiture 1 presents the construction blocks of tag graph. The basic elements are $TagNode$ and $TagEdge$, they are $subclasses$ of $TagElement$ and constitute the $TagGraph$. We develop the idea of graphic tagging to embody more complicated and detailed information following the RDF[4] graph model with respect to its merits: the capability of "say anything about anything", distributable, scalable and easy to connect each other.

The concept $Context$ in Figure 1 connects the tagger (modelled as $FOAF : person$), the tag ($TagGraph$) and the web resource be tagged. $Context$ catches the scene of the tagging action, hence the tagging semantics can be induced from the scene instead of from keywords only. Currently, the tagged object is presented by a URI for the implementation

[4] `http://www.w3.org/RDF/`

simplicity. In the future more complicated structures should be considered and can be plugged in.

A tag might links to its object (e.g. web page) in one of many implicit meanings as discussed in Section 2.1. In this work we adopt the "linking" semantics defined in [Bechhofer *et al.*, 2002] : The tag graph just *appear* in a *context* instead of belonging to the web content. This decision is made by our attempt to decouple the tag and the web content. We argue that the tag graph in the end, should be the semantic equivalent to its tagged object, i.e., a machine understandable counterpart of the human readable content.

The last argument is that the word used for tagging should be viewed as neither an ontology concept nor an instance. It is just a symbol whose meaning changes according to the user and the circumstance. Thus we provide the $stickTo$ property which is used to define the semantics of a tag element dynamically. By "sticking" a tag element "to" the resource in a knowledge base, its meaning can be assigned (and changed when necessary) without imposing any controlled vocabulary on the user. At the same time machine processing ability can be achieved as well.

To summarize, in this work the tag ontology is designed to support the emergent semantic. Based on the tag ontology, a graph structure instead of a bunch of keywords is used to describe (tag) the web content. The scene for tagging activity is recorded which can be exploited in the future to approximate the tagging semantics as preciously as possible. And the semantics of a tag element can be defined in a formal, dynamic fashion.

2.3 Layered Semantics

As discussed before, folksonomy system provides services that enable the user to consume metadata he endeavored to create immediately. The immediate benefits is core to the success of folksonomy and is emphasized in our design of tag ontology as well.

Briefly speaking, the tag ontology embodies a semantic layer consisting of three compatible levels that support various semantic operations, varying from simple to complex. Autonomous agents can commit to and work at any of the three levels. Thus the relatively complex model (tag ontology) not only does not lose the "immediate benefits", but provides the basis for more services. The details of the three semantic levels is listed as below:

- Keyword level. At this level, the tag is keyword attached to the web content, and the tag graph degenerates into a set of tag nodes without edges. Existing folksonomy system works at this level.
- Tag graph level. Here, the tag graph formed by the triple statements tags the web content as a whole. Distributed tag graphs can be connected and combined. Applications working at this level commit to all the terms from the tag ontology except the $stickTo$ property.
- Semantic level. This level is compatible to the tag graph level in that tags are viewed as triple statements too. However, with $stickTo$ property, A $TagElement$ can point to some knowledge base and possess the formal semantics.

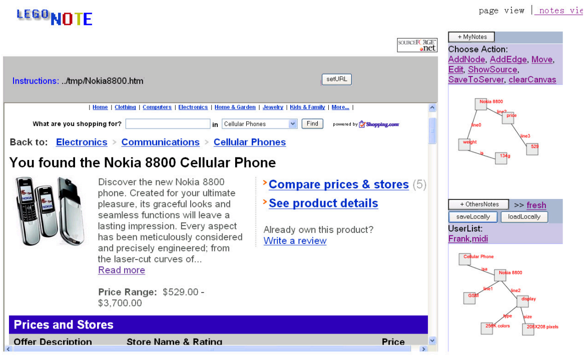

Figure 2: How a user tags a page with Lego-Note. A web page about a cellular phone is loaded in the embedded browser. The user's tag graph is shown in the upper right area. The same page is also tagged by two other users: Frank and midi, Frank's graph is also loaded and shown.

In this sense, the tag ontology serves two functions. First, the social metadata organized around tag ontology can be used by the Semantic Web community directly as semantic metadata. Second, the tag ontology structurally supports the emergent semantics, and offers a seamless incremental upgrade path for emergent semantics and semantic services without requiring a wholesale adoption of the Semantic Web's vision.

3 Implementation

To explore the idea of graphic tagging an experimental system, Lego-Note is implemented. The implementation detail will be introduced in this section.

3.1 User Experience

First, we introduce the functions implemented in Lego-Note[5] by how a user might experience it.

Figure 2 presents the user interface for the user to tag a page. The left-hand side of Figure 2 is an embedded browser, which consumes most of the space in order to preserve the familiar browsing experience. The user can browse web pages by either entering a URL to the embedded browser or following the links in the web page.

The graphic tagging tool is floating on the right-hand side. In the upper area, the user can create and edit his own tag graph. Basic graphic operations are provided interactively, showing themselves as links ("AddNode", "AddEdge", "Move" and etc.) in the toolbar . The graph as a whole tags the web content to its left. A graph node describes either the whole page or the user selected page fragment. In the latter case, the node saves the position and content of the fragment so that it can be relocated next time. The user can share his tag graph by uploading to the server. The bottom area exhibits tag graphs to the same page shared by the others. Others' tag

[5]The online demo can be accessed from `http://dom-sensus.sourceforge.net/`

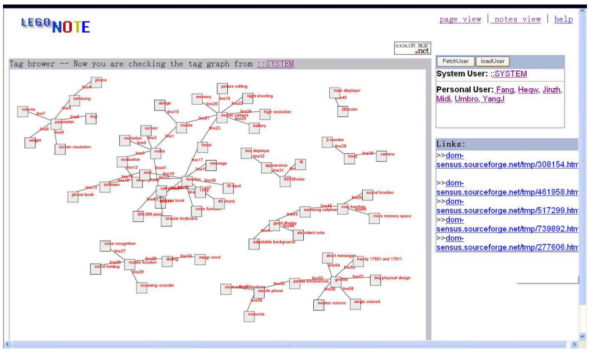

Figure 3: The tag browser. It organizes information around tags.

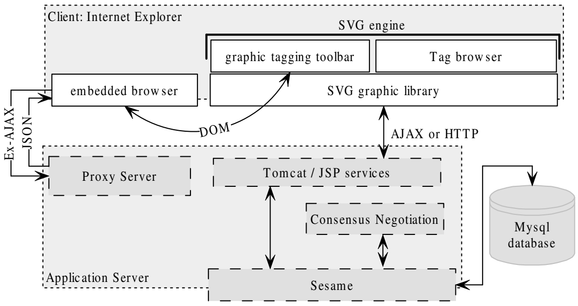

Figure 4: Lego-Note is an AJAX application taking the browser-server architecture.

graphs help the user catch the important information more quickly and build his own more easily.

In Figure 2, a web page about the cellular phone is loaded and the tag graph of the current user is shown in the upper right area. The other two users, namely, Frank and Midi have shared their tag graphs. Currently Frank's graph is exhibited. Figure 2 contains the functions organized around a web page. It is called a page view of Lego-Note. Also a tag view is provided as well, where the information is organized around tags. The tag browser that provides the tag view is depicted in Figure 3.

There are three function areas shown in Figure 3. The right-upper area lists the names of system users. Click the user name, and the summarized tag graph will be presented in the central (left) area. The summarized graph is generated by the system by linking all the tag graphs of a user. Click a node of the summarized graph, URLs being tagged by the node will be listed in the left-bottom area. There exists a special user ::SYSTEM whose tag graph is generated from that of every system user. Figure 3 shows the tag view of the user ::SYSTEM.

3.2 Implementation

Lego-Note is an AJAX (Asynchronous JavaScript And XML) application based on the Browser-Server structure. Figure 4 gives its overall architecture.

On the client side, an SVG (Scalable Vector Graphics[6]) graphic library is implemented and the floating toolbar for graphic tagging and the tag browser are built based on it. The graphic tagging toolbar interacts with the embedded browser by DOM[7] operations. All the HTTP requests initiated from the embedded browser are sent to a proxy server, instead of to the web directly. The proxy server then fetches the page from the web and sends the returned page content back to the embedded browser in the form of JSON (JavaScript Object Notation). Communications between the embedded browser and the proxy sever are carried out using EX-AJAX[8]. This additional process is necessary to bypass the "same domain" security restriction. Otherwise, the tagging toolbar is forbade to manipulate the DOM of the web page.

On the server side, Tomcat application server runs to process the user requests, and Sesame[9] manages the storage, querying and inferencing of the tag graph which is modelled by the tag ontology. Lego-Note works on the "tag graph level" of the semantics layer discussed in section 2.1

Although Lego-Note is implemented with the hope of environment independent, the various implementations of the web browser (e.g., firefox and Internet Explorer) and SVG viewers makes it a time-consuming work. At present Lego-Note only works with International Explore version 5.0 above and Adobe svg viewer version 3.0 above.

4 Case Study

In this section, we provide a preliminary case study and discuss how and to what extent the consensus might be achieved based on the tag graphs from different users with the Lego-Note as the experimental environment.

4.1 Experimental Setup

First, we prepared a data set by downloading 300 web pages from `http://www.pconline.com.cn/mobile/review/`, a web site where users can write their reviews about the mobile phone. Then a video clip is made, demonstrating how to use the system by creating the tag graph of a researcher's home page. The web content demonstrated in the video is in a different domain other than the mobile phone because we tried to avoid biasing the testers. At last, 4 college students are invited to the experiment. After learned how to use system from the video, they are asked to tag web pages and to write down their comments. The web pages are selected randomly from the dataset.

4.2 Result and Analysis

There are 11 tag graphs constructed for 11 different web pages by the four users. The detailed data is presented in Table 1. Take user "A" as an example, he tagged two pages, 16 nodes and 13 edges are created in all of his tag graphs. The number of words (separated by blank space) he used for graph nodes is 29 and that for graph edge is 16. For all the figures listed in Table 1, repeated elements (e.g. the node

[6] `http://www.w3.org/TR/SVG/`
[7] `http://www.w3.org/DOM`
[8] `http://ajaxextended.com/`
[9] `http://www.openrdf.org/about.jsp`

Table 1: Tagging information summarized from 4 users on 11 web pages.

	A	B	C	D	SUM
Page No.	2	2	4	3	11
Node No.	16	19	29	40	104
Edge No.	13	18	26	31	88
Word No. of Node	29	27	37	109	202
Word No. of Edge	16	18	26	52	112

name, the edge name or the word) are counted separately. The last column lists the sum of all users. The numbers for node and "words of node" are listed separately because the user might use several words in a single node, e.g. the node "$soundfunction$" contains two words. The same reason applies to the numbers for edge and "words of edge".

The analysis and observations are given in the following part of this section.

Users tend to get words connected. From Table 1 we can see that the number of edge is close to that of the node, which shows optimistically that people are willing to get words related when lead by convenient tools. However, although edges are added to the tag graph frequently, they are left unnamed quite often. In the experiment, numbers of unnamed edge are 6, 18, 26 and 0 for each user respectively. Generally users leave some edges unnamed with user B and C naming none. This might imply that users are most interested in the information "directly" related to the content (the web page). Because edges describe relations among nodes (keywords) which is a step further from the content, they don't appeal the users so much.

We then connect the 11 tag graphs by the nodes holding the same names thus form the summarized tag graph, as is shown in Figure 6. Compared with the ontology constructed from classic folksonomy system given by [Mika, 2005], the summarized tag graph reveals much more details of the domain that possesses a mixture of conceptual words (e.g. mobile phone, camera, good and etc.) and instantial words (e,g, 130g, 260 color, 2-monitor and etc). In other words, the summarized tag graph looks more like an instantiated ontology than an ontology with defined vocabulary only. This feature is brought by the adoption of the graph structure that allows the user makes detailed statements into the web content.

Vocabulary consensus can be summarized. The consensus on vocabulary varies with participants and the depth of agreements. In this experiment we exam the popularity of words first. Figure 5 presents the distributions of five variables summarized from all tag graphs. The five variables are: the node, the edge, the word used in node, edge and both of two. As can be seen from Figure 5, the distribution curves show similar properties as that from [Michlmayr, 2005] and [Shaw, 2005]: Popular words decreases very rapidly and the resultant curve falls asymptotically towards y=1. It supports the view of "there is a natural tendency towards the convergence of tags and that strategies to facilitate this development exist"[Guy and Tonkin, 2006].

Structure consensus need to be explored further. Except the vocabulary consensus we would like to discuss the feature of structure consensus, which represents the users' agreements on the possible relationships among keywords. The assumption is, with the graph tagging tool users are expected to connect tag nodes more frequently, hence relations among vocabularies are more likely to be exposed.

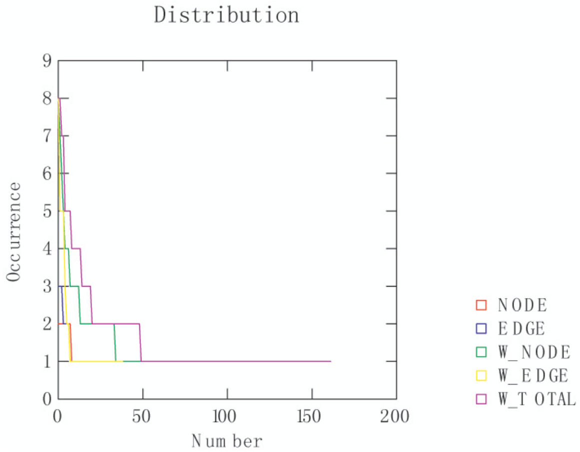

Figure 5: The distributions from tag graphs. Five curves show the number distributions of the node, edges, word used for node, edge and for both respectively.

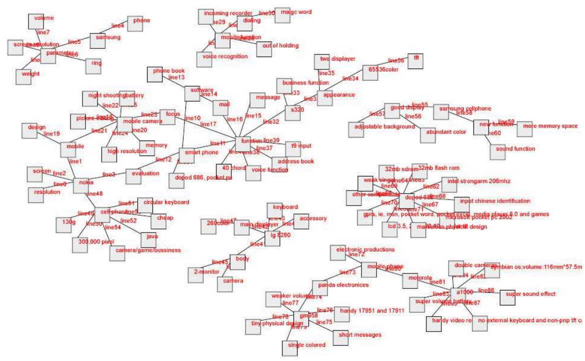

Figure 6: The summarized tag graph from all the users

As discussed above the user tends to get nodes connected when creating his own tag graph. However, there are not so many interconnections among graphs showing in Figure 6. Most of the time the tag graph of a web page stands alone by itself and occasionally connects to each other by the nodes possessing the same name. Further more, no structure overlapping (overlapped triple statements) occurs.

The shallow structure of the tag graph mainly answers for the lack of structure consensus. Theoretically with the tag graph, a user can make any statement in the triple form, while in fact almost no one shows interests in explaining something out the context of the web page shown to him. Except user D, who adds a triple claiming that $< mobilephones, isa, ElectronicProductions >$. The observation proves again that the immediate benefits should be provided in order to persuade social users to contribute metadata. Besides, we can deduce that users play different roles on how they contribute to the community vocabulary.

The proposed scenario. Except the tag graphs constructed during the experiment, we also collected users' comments about graphic tagging. Two questions are posed by all of the users. The first one is that the graphic tagging tool is interesting but not convenient, the second follows as "what's the benefit that comes from these tag graphs".

The first question reveals the appeal for the mechanism that extracts triples from phrases (semi-)automatically. An other observation also supports the demanding: to describe some thing in detail instead of endeavoring to build triples, the user tends to use phrases as the name of a node that can be split into triples in-depth. A better choice might be to let the user enter simple phrases that can be transformed into tag graph (semi-)automatically, or to extract triples from the user selected contents directly.

To argue its usefulness, consider the scenario of online auction or online flea market like Yahoo! auction and eBay. Typically, to post an item for sell the seller is required to fill in a form describing the attributes of the item and to write down a free-text description as well. On the other hand, the buyers can search for their favorite items based on predefined features. These services are typically based on the commodity templates. However, the shortcoming also rises from the templates, which try to covering every aspects of commodities and usually become too general. Besides, the template based implementation is hard to keep itself up to date. By providing a flexible infrastructure based on the tag ontology, which allows users to post their own statements then generate the templates from the social metadata dynamically, we might conquer the problem. A system that embodies the similar idea is google base[10].

5 Conclusion and Future Plan

Nowadays, the in-depth study and wide adoption of Semantic Web meet a barrier since there lacks semantic metadata. The top-down approach of semantic metadata creation by means of ontology based annotation is insufficient in ad-hoc and dynamic situations. Folksonomy which is the result of shared tags created by normal web users is drawing attention as a promising source of semantic metadata and the emergent semantics is proposed to support the machine processability. In this paper we propose a concept model that supports metadata generation by extending the ideas from folksonomy. The layered semantics is proposed in the model which provides a gradual way to generate semantic metadata. An implementation system, Lego-Note, based on the model is introduced, which accumulates metadata in a social and bottom-up approach. With Lego-Note users can create tag graph about web pages freely. To shape the domain consensus from social tag graphs, a primary experiment is performed, it servers as a case study providing valuable clues for the future work.

The future work can be considered from three aspects. The first one relates to generate and purify the tag graph, including for example, eliciting triples from phrases. The second aspect concerns how to achieve consensus that varies with community and time. The last one, as mentioned in Section 4, (semi-)automatical mechanisms should be introduced into Lego-Note, extracting triple from phrases.

[10] `http://base.google.com/`

References

[Bechhofer *et al.*, 2002] Sean Bechhofer, Les Carr, Carole A. Goble, Simon Kampa, and Timothy Miles-Board. The semantics of semantic annotation. In *On the Move to Meaningful Internet Systems, 2002 - DOA/CoopIS/ODBASE 2002 Confederated International Conferences DOA, CoopIS and ODBASE 2002*, pages 1152–1167, London, UK, 2002. Springer-Verlag.

[Berners-Lee *et al.*, 2001] T. Berners-Lee, J. Hendler, and O. Lassila. The semantic web. *Scientific American*, 284(5):35–43, 2001.
Essay about the possibilities of the semantic web.

[Cudre-Mauroux *et al.*, 2006] Phillippe Cudre-Mauroux, Karl Aberer, Alia I. Abdelmoty, Tiziana Catarci, Ernesto Damiani, Arantxa Illaramendi, Mustafa Jarrar, Robert Meersman, Erich J. Neuhold, Chrisine Parent, Kai-Uwe Sattler, Monica Scannapieco, Stefano Spaccapietra, Peter Spyns, and Guy De Tre. Viewpoints on emergent semantics. *Journal on Data Semantics VI*, pages 1–27, 2006.

[Domingue and Dzbor, 2004] John Domingue and Martin Dzbor. Magpie: supporting browsing and navigation on the semantic web. In *Proceedings of the 9th international conference on Intelligent user interface*, pages 191–197. ACM Press, 2004.

[Fensel *et al.*, 2002] D. Fensel, J. Hendler, H. Lieberman, and W. Wahlster. *Weaving the Semantic Web*. MIT Press, 2002.

[Gruber, 2005] Thomas Gruber. Ontology of folksonomy: A mash-up of apples and oranges. Technical report, November 2005.

[Guy and Tonkin, 2006] Marieke Guy and Emma Tonkin. Folksonomies: Tidying up tags? *D-Lib Magazine*, 12(1), January 2006.

[Handschuh *et al.*, 2001] Siegfried Handschuh, Steffen Staab, and Alexander Maedche. Cream: creating relational metadata with a component-based, ontology-driven annotation framework. In *K-CAP '01: Proceedings of the 1st international conference on Knowledge capture*, pages 76–83. ACM Press, 2001.

[Huynh *et al.*, 2005] David Huynh, Stefano Mazzocchi, and David Karger. Piggy bank: Experience the semantic web inside your web browser. In *Proceedings of the 4th International Semantic Web Conference (ISWC)*, 2005.

[Kahan and Koivunen, 2001] José Kahan and Marja-Ritta Koivunen. Annotea: an open rdf infrastructure for shared web annotations. In *WWW '01: Proceedings of the 10th international conference on World Wide Web*, pages 623–632, New York, NY, USA, 2001. ACM Press.

[Karl *et al.*, 2004] ABERER Karl, CUDRE-MAUROUX Philippe, OUKSEL Aris M., CATARCI Tiziana, HACID Mohand-Said, ILLARRAMENDI Arantza, KASHYAP Vipul, MECELLA Massimo, MENA Eduardo, NEUHOLD Erich J., DE TROYER Olga, RISSE Thomas, SCANNAPIECO Monica, SALTOR Flix, DE SANTIS Luca, SPACCAPIETRA Stefano, STAAB Steffen, STUDER Rudi, YoonJoon Lee, Jianzhong Li, Kyu-Young Whang, and Doheon Lee. Emergent semantics principles and issues. In *Database Systems for Advanced Applications 9th International Conference, DASFAA*, pages 25–38. Springer, March 2004.

[Kettler *et al.*, 2005] Brian Kettler, James Starz, William Miller, and Peter Haglich. A template-based markup tool for semantic web content. In *Proceedings of the 4th International Semantic Web Conference (ISWC)*, 2005.

[Kiryakov *et al.*, 2003] Atanas Kiryakov, Borislav Popov, Damyan Ognyanoff, Dimitar Manov, Angel Kirilov, and Miroslav Goranov. Semantic annotation, indexing, and retrieval. In *Proceedings of the 2nd International Semantic Web Conference*, pages 484–499, 2003.

[Michlmayr, 2005] Elke Michlmayr. A case study on emergent semantics in communities. In *Workshop on Semantic Network Analysis, International Semantic Web Conference (ISWC2005)*, 2005.

[Mika, 2005] Peter Mika. Ontologies are us: A unified model of social networks and semantics. In *Proceedings of the 4th International Semantic Web Conference (ISWC 2005)*, 2005.

[Shaw, 2005] Blake Shaw. Building a better folksonomy: Web-based aggregation of metadata. Technical report, 2005.

[Uschold, 2003] Michael Uschold. Where are the semantics in the semantic web. *AI Magazine*, 24(3):25–36, 2003.

[Wu *et al.*, 2005] Xian Wu, Lei Zhang, , and Yong Yu. Exploring social annotations for the semantic web. In *Proceedings of the 15th international Conference on World Wide Web (WWW '06)*. ACM Press, 2005.

[Xu *et al.*, 2006] Zhichen Xu, Yun Fu, Jianchang Mao, and Difu Su. Towards the semantic web: Collaborative tag suggestions. In *Proceedings of the Collaborative Web Tagging Workshop at the WWW 2006*, 2006.