

Between Ontology and Folksonomy: A Study of Collaborative and Implicit Ontology Evolution

Jiahui Liu

Northwestern University
j-liu2@northwestern.edu

Daniel M. Gruen

IBM Research
daniel_gruen@us.ibm.com

ABSTRACT

We present our first user study of CRAFT, a semantic prototype for investigation and analysis in which users collaboratively and implicitly extend the system's ontology to capture new concepts as they conduct their work. We devised a paradigm in which multiple series of ontologies evolve in different trajectories from the same initial point. We analyze the ontology evolution quantitatively with several metrics, and user behavior qualitatively through interviews and observation. Based on our study, we propose a set of design suggestions for semantic applications with collaborative and implicit ontology development.

Author Keywords

Ontology, Collaboration, Folksonomy, User Study

ACM Classification Keywords

H.5.2 GUI; Evaluation/methodology, H.5.3 CSCW

INTRODUCTION

Ontologies are a core technology for knowledge management and semantic web applications. Traditionally, ontology development is approached in a centralized and formal way. Expert knowledge engineers are responsible for creating and maintaining ontologies for separate larger communities of end users of the semantic applications. The centralized approach to knowledge management leads to a number of problems, making ontologies expensive to build and hard to keep up-to-date with changes in the domain and user needs. Moreover, lack of communication between creators and users incurs misunderstanding and cognitive constraints in using ontologies [3].

In contrast to the centralized formal approach in ontology engineering, social web applications (e.g. Flickr and del.icio.us) have shown great success in encouraging common users to create, tag, and share content. Informal knowledge emerges in the form of *folksonomy* as users are performing their daily tasks, such as sharing pictures and

organizing bookmarks.

Motivated by the challenges of centralized ontology engineering and inspired by the success of social web applications, researchers have proposed conceptual frameworks and tools that leverage collaborative and implicit knowledge construction and sharing for formal knowledge acquisition in semantic applications [1, 4]. In this new approach, knowledge engineering is implicitly integrated with the work process of end users. The community collaboratively shapes the formal domain definitions through sharing and negotiation.

Collaborative and implicit ontology evolution has the potential to bridge the gap between knowledge engineering and usage. However, there has been little research into how non knowledge engineering experts will approach ontology engineering in their daily work. It leaves open the question how the ontology will change in the process. In this paper, we present a user study that investigates users' behavior and ontology evolution with a prototype system called CRAFT (Collaborative Reasoning and Analysis Framework and Toolkit). CRAFT assists groups of analysts, such as market researchers, in collecting information, making inquiries and reaching decisions. In CRAFT, users are empowered to extend the ontology to capture new concepts on demand as they are conducting their investigations. In our user study, we explored the following research questions: 1) *Are users able to create and maintain ontologies while they are engaged in their knowledge-intensive work?* 2) *How will the ontology evolve in the hands of users without knowledge engineering experience?* 3) *How similar are the ontologies created about the same domain by different users?* 4) *How will the design of the tool affect users' behavior and ontology evolution?*

We believe our questions and findings will give fresh insight into collaborative and implicit ontology evolution in semantic applications. Based on our user study, we propose a set of design suggestions for CRAFT and similar systems.

CRAFT

CRAFT supports the sensemaking activities of analysts, including collecting information, making inquiries, analyzing results and sharing findings. Within CRAFT, users can frame their problem and express their understandings with graphical models. CRAFT uses ontologies to capture domain semantics and enable sharing

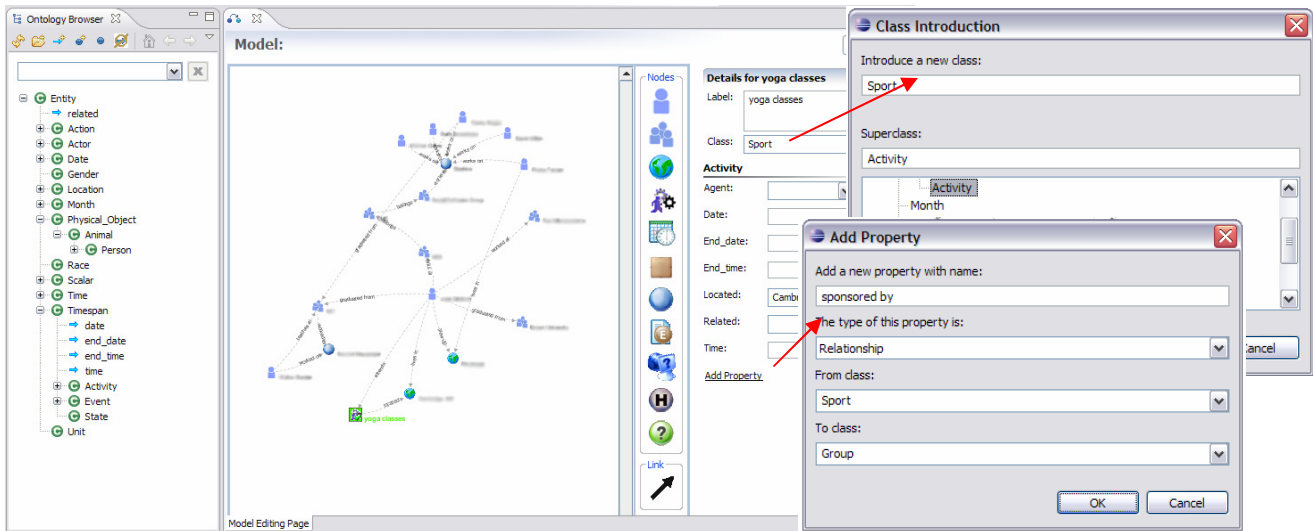


Figure 1. Introducing new class and relationship on demand in CRAFT

among users and with external systems. CRAFT users are able to extend the ontology on the fly as they uncover and record new concepts in their investigation.

Figure 1 shows a screenshot of CRAFT. In the center is the *model editor*, where analysts conduct their investigations. Users create their models by dragging icons from the palette into the model and connecting them with links. Each icon corresponds to a commonly used class in the ontology, shown in the *ontology browser* on the left. Users may search the ontology and drag classes or entities from the ontology into their models. When a user selects an element in the model, the *details pane* on the right will show the semantic information about the element, specifically, class and properties for entities, and labels for relationships. The user may edit the information, assisted by auto completion. If a class, relationship, or property input by the user is not in the ontology, CRAFT will prompt the user to add it.

USER STUDY

We conducted a study to analyze the ontology evolution when a succession of users implicitly extended an ontology in their work process. The paradigm we devised allows for multiple evolution trajectories from the same initial point to be observed, allowing for comparison and analysis.

Method

We recruited nine interns in our research group (referred to as S1-S9) to use CRAFT for an investigation task. None of the subjects had previous experience in knowledge engineering, nor had they used CRAFT before.

Subjects were given a 15 minute introduction to the basic functionalities of CRAFT. The subjects were then given 30 minutes to collect information about a colleague researcher and create a graphical model to record the information. We chose a topic familiar to the subjects to alleviate the cognitive load of comprehension, so subjects could focus on the investigation. We provided several web pages about

basic ontology	Series 1	Series 2	Series 3
1 st generation	S1 model R1	S2 model R2	S3 model R3
2 nd generation	S4 model R3	S5 model R1	S6 model R2
3 rd generation	S7 model R2	S8 model R3	S9 model R1

Table 1. User study paradigm

the researcher as a starting point. Subjects were also free to search online for additional information as they wished.

To analyze the evolution of the ontology in collaborative use, we divided the nine subjects into three groups. All groups started with the same impoverished ontology. Within each group, the subjects took turns creating a model, extending the ontology as needed. Table 1 illustrates the arrangement. For example, S1 created a model about R1 with the basic ontology and extended the basic ontology to the 1st generation. Afterwards, S4 created a model about another researcher R2 with the 1st generation ontology, resulting in the 2nd generation. As shown in Table 1 each group created models for the same set of researchers, but in different order. Thus, we obtained three ontology series evolving on different paths from the same starting point.

It should be noted that we did not specifically bias the subjects to being careful in extending the ontology. The subjects were told “*feel free to use anything existing and create anything necessary. The changes you made to the ontology will be available to other interns to use*”. We conducted interviews with the subjects after they completed their investigations, to further understand their experience and the strengths and weaknesses of the UI.

ANALYSIS AND FINDINGS

Model creation

In 30 minutes, the subjects created fairly complicated models, containing 13 to 26 nodes and 16 to 26 links, (averaging at 17.7 nodes and 20.9 links per model). A followup study showed that others were able to understand

the models and extract information from them. Given the fact that this is the first time they used CRAFT, the result is encouraging. It shows that the users are able to effectively express their ideas using the tool. In the interview, all of the subjects said using the ontology was not constraining. As said by S7, “*I can always create something new if it is not there*”. Thus, integrating ontology extension seamlessly with end user tasks is a promising approach to addressing constraints of ontology enabled semantic applications.

Ontology evolution

As we expected, subjects created new classes and relationships during their investigations. Consequently, the size of ontologies increased, as illustrated in Figure 2. It is evident in the figure that subjects created relatively few new classes. Closer investigation reveals that classes used by the subjects show characteristics of “basic level” categories. Reflecting the cognitive aspect of taxonomies of common objects, the basic level (such as “chair”) is most informative, as opposed to superordinate levels (such as “furniture”) and subordinate levels (such as “wood armchair”) [6]. Experiments demonstrate that the basic levels are most quickly identified and generally agreed upon [6]. When subjects categorize the entities in their investigation, they would most likely categorize them into the basic levels, like “Person” and “Company”. A similar phenomenon is reported in social tagging of bookmarks [2].

We further analyzed the change of the ontology structure in terms of *inheritance richness (IR)* and *relationship richness (RR)*, two metrics proposed by Tartir et al. [7] for analyzing ontology schemes. Figures 3 and 4 plot the IR and RR respectively for all generations in the three ontology series.

IR is the average number of subclasses per class [7]. High IR indicates horizontal ontology with shallow specialization, and vice versa. Figure 3 shows increases of IR in ontology evolution for all three series. As discussed above, subjects tended to create “basic level” classes. They did not bother to create abstract superclasses or further specify with subclasses. Consequently, the ontologies become flatter as they evolved.

Figure 4 demonstrates that the ontologies became richer with relationships between the classes. In contrast to the classes, the subjects tried to make the relationships explicit. As S6 said, “*I wanted to put more information to the links*”. This is evident in the names of the relationships, which usually consisted of more than two words. The subjects tended to make a whole sentence with the two entities and the relationships connecting them. The subjects were seen reversing the direction of the links to make sure it “*sounds right*” (S9). Similar user behavior was observed in the study of ClaimSpotter, a semantic annotation tool [8].

Similarity between ontology series

We now turn to our third research question: *How similar are the ontologies created about the same domain by different users?* We examine this question by comparing the

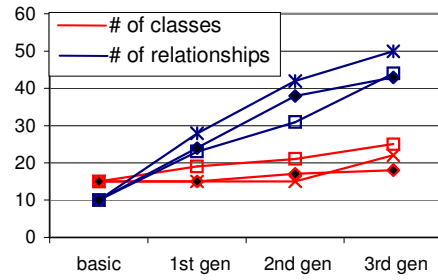


Figure 2. Change of ontology size

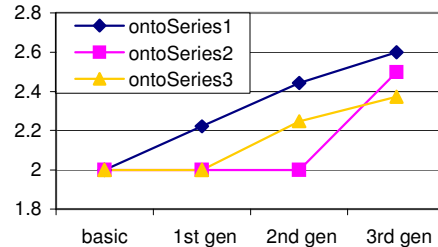


Figure 3. Inheritance richness of ontologies

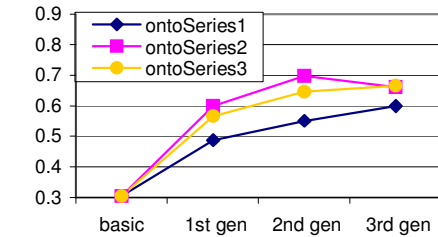


Figure 4. Relationship richness of ontologies

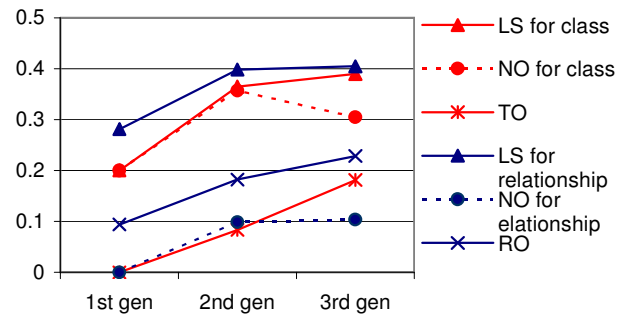


Figure 5. Similarity of ontologies

ontology series using the ontology similarity measures proposed by Maedche and Staab [5]. The analysis can provide insights into mechanisms for recommending similar concepts to users to foster knowledge reuse.

For comparison at the lexical level, Maedche and Staab proposed the *lexical similarity (LS)* measure [5], which diminishes the variation of word forms, such as singular vs. plural. We supplemented this matrix with the *name overlap (NO)* to measure the overlap of exact names. At the conceptual level, Maedche and Staab proposed two metrics [5], *taxonomic overlap (TO)*, for the class-superclass semantic structure, and *relation overlap (RO)*, for the similarity of relationships in terms of their domain and range classes. In our analysis, we extracted the classes and relationships not included in the starting basic ontology, i.e. the new concepts added by the subjects. We then applied the comparison measures to the sets of new concepts of

each generation in the three ontology series. Figure 5 illustrates the average similarity of each generation.

Figure 5 shows that the three ontology series about the same domain are different from each other. However, as more subjects added more concepts, the similarity increased over generations. An interesting phenomenon to note in Figure 5 is that the LS and NO for class are very close, but not for relationships. According to our observation, concepts with high LS in our experiments are mostly similar concepts. The figure reveals that subjects have high agreement on names of similar classes, but low for similar relationships. As discussed above, names of relationships are more explicit and longer, so they are prone to more variation. Another reason for lower agreement relates to the direction different subjects chose for basically the same relationship (e.g. “employer of” vs. “employed by”).

DESIGN IMPLICATIONS

The user study suggests a number of design implications for our system and others with similar goals.

Scaffolding for ontology development

With collaborative implicit ontology evolution, the users who modify the ontologies are not expert knowledge engineers. Furthermore, their main goal is to accomplish their tasks, rather than to create an effective ontology. Based on our study, we suggest two kinds of scaffolding to improve resulting ontologies. Firstly, the ontology view should be integrated within the interface for introducing new concepts, so the users will be aware of the taxonomic structure when they are extending the ontology. Secondly, the system should provide descriptions and examples to help users understand the concepts in the ontology, such as representative subclasses and frequently used entities for classes, especially abstract classes.

Incremental formalization

The ontologies that resulted in our study show a combination of aspects of ontologies and folksonomies. Although the concepts were organized into the taxonomy, certain abstraction levels were missing. Subjects were sometimes aware of this: for example, S9 explicitly reported that he “*almost thought of creating something like that (an abstract class) if there were enough classes requiring this superclass*”. We propose investigating mechanisms by which the system could detect situations that would benefit from consolidation of classes or other ontology editing, and prompt the user to do so. The system could facilitate collaboration among people using the relevant concepts in making such decisions, and potentially involve expert knowledge engineers as well.

Searching vs. auto completion

CRAFT includes search and auto completion features to encourage reuse of concepts already in the ontology. Subjects seldom used the search functionality in our study, but often accepted the suggested autocompletions. As S6 said in the interview, “*I would only search if I thought it was there*”. Our analysis of the similarity between ontology

series suggests that a more sophisticated auto completion be implemented which would handle variations of concepts’ names and structural differences. Specifically, in recommending relationships, similar relationships with different directions should also be detected. We also have added a mechanism for user specification of synonyms.

CONCLUSION AND FUTURE WORK

In this paper, we describe a paradigm for investigating the evolution of ontologies as non expert users extend ontologies in their work process. Our preliminary study suggests that seamlessly integrating ontology engineering with users’ tasks is a promising approach, showing a combination of aspects of ontology and folksonomy. Based on our study, we propose design suggestions for semantic applications with collaborative and implicit ontology engineering. Note that in our study, automatic information gathering was not part of the task. Subjects could extend the ontologies at will to express their ideas, without considering how their changes would affect integration with external systems. In our future work, we will investigate the extent to which such integration increases motivation to conform to existing concepts. We also plan additional testing with more users and ontology iterations with this paradigm, and will use it to evaluate the design suggestions we propose.

ACKNOWLEDGEMENTS

We thank Steven Ross, Susanne Hupfer, John Patterson and Jamie Rasmussen for their work to build the system described here, developed as part of the IBM Research System S project.

REFERENCES

1. Braun, S., Schmidt, A., Walter, A.. “Ontology Maturing – A collaborative Web 2.0 Approach to Ontology Engineering”, *Workshop on Social and Collaborative Construction of Structured Knowledge at WWW*, 2007.
2. Golder, S. A., “Usage patterns of collaborative tagging systems”, *Journal of Information Science*, 32(2), 2006.
3. Hepp, M. “Possible ontologies: How reality constraints building relevant ontologies”. *IEEE Internet Computing*, 11(1), 2007.
4. Kotis, K., Vouros, G. A., and Alonso, J. P.. “HCOME: A Tool-Supported Methodology for Engineering Living Ontologies”. *Semantic Web and Database Workshop* 2004.
5. Maedche, A., Staab, S., “Measuring Similarity between Ontologies”, *International Conference on Knowledge Engineering and Knowledge Management. Ontologies and the Semantic Web*, 2002.
6. Tanaka, J., and Taylor, M., “Object Categories and Expertise: Is the Basic Level in the Eye of the Beholder?” *Cognitive Psychology* 23(3), 2005.
7. Tartir, S., Arpinar, I. B., Moore, M., Sheth, Aleman-Meza, A., B., “OntoQA: Metric-Based Ontology Quality Analysis”, *IEEE ICDM Workshop on Knowledge Acquisition from Distributed, Autonomous, Semantically Heterogeneous Data and Knowledge Sources*, 2005.
8. Sereno, B., Buckingham Shum, S., and Motta, E., “Formalization, User Strategy and Interaction Design: Users’ Behaviour with Discourse Tagging Semantics”. *Workshop on Social and Collaborative Construction of Structured Knowledge at WWW*, 2007.