Collaborative Reasoning and
Collaborative Ontology Development in CRAFT

Daniel Gruen and Jamie Rasmussen
Collaborative User Experience Group
IBM Watson Research Center
Cambridge, MA 02142
{daniel_gruen,jrasmus}@us.ibm.com

Jiahui Liu
Intelligent Information Laboratory
Northwestern University
Evanston, IL 60208
j-liu2@northwestern.edu

Susanne Hupfer and Steven Ross
Collaborative User Experience Group
IBM Watson Research Center
Cambridge, MA 02142
{shupfer,steven_ross}@us.ibm.com

Abstract

We present CRAFT (Collaborative Reasoning and Analysis Framework and Toolkit), a tool for collaborative investigation, reasoning, and analysis. Analysts use CRAFT to represent their collective knowledge and reasoning via interconnected graphical models built upon a shared evolving ontology. These semantic models help connect analysts to digital information sources and to each other, and the aggregated knowledge and findings of many analysts may be analyzed and visualized. We also summarize the results of a preliminary user study of collaborative, implicit ontology evolution using this tool.

Introduction

There are many situations in which an organization or group of people must work together to collect information and reach a consensus on what is happening. There is often too much information for one person to sift through alone, and the required expertise may be spread among many individuals. In fields such as business intelligence, risk analysis, fraud detection, homeland security, financial forecasting, epidemiology, and strategic planning, there is a need for groups of people to share information and reason together, and our work is focused on helping them to do so more effectively. The Collaborative Reasoning and Analysis Framework and Toolkit (CRAFT) is a research prototype aimed at helping analysts as they collect and share information to support decisions. CRAFT lets analysts represent what they know about a situation, record questions and hypotheses, and create inquiries for new information from internal databases and public sources – using concepts and instances drawn from an evolving ontology. CRAFT includes facilities to keep analysts aware of new information and inquiry results, reuse the information added by other analysts, and collaborate across investigations and roles.

The philosophy underlying the CRAFT approach was influenced by wiki software, in which users can easily add, edit, and link information on various topics, continually

refine the topic descriptions, and compare them with previous versions. Unlike typical wikis, however, information in CRAFT is semantically encoded and added through a combination of graphical and form inputs. The targeted users of our system are knowledge workers but not knowledge engineers, so we aim for a knowledge representation language that balances expressivity and ease of use.

In contrast to traditional centralized approaches to ontology development, which can result in ontologies that are hard to maintain and constraining to use, CRAFT empowers users to extend the ontology on demand, by capturing new classes and properties as they conduct their investigations. Changes made by each user are immediately available for others to use, resulting in a community resource that organically grows and adapts to its users' needs.

CRAFT is designed to support an investigation life cycle, in which analysts pose questions, gather and organize evidence, evaluate hypotheses, and are led to new questions by the resulting conclusions. CRAFT also serves as the analyst's interface to System S, a large-scale distributed stream processing system being developed in IBM Research.

This paper is structured as follows: We first discuss knowledge representation and modeling features in CRAFT, including the ways a user adds information to the system and extends the ontology. We describe some of the expressivity/usability tradeoffs we have made, and describe several user interface components for browsing the shared ontology. Next, we discuss the role inquiries play in gathering information and evidence, and how inquiries in CRAFT differ from web searches and database queries. We then describe some of the external software packages we have embedded to provide analysis and visualization of inquiry results and aggregated claims. Finally, we discuss results from a preliminary user study of ontology evolution in CRAFT.

Knowledge Representation and Modeling

In CRAFT, each analyst can belong to multiple ongoing investigations. Each investigation provides a kind of

private team space for members to work together to achieve a goal, answer a question, or research a problem or situation. An investigation can contain any number of models. Each model contains references to entities (instances of classes) and claims about those entities. Each entity may optionally have one or more "primary" models, which are models about that entity. For example, a competitive analysis model might reference an IBM entity of class Company, and that entity may be linked to several models of the IBM Corporation. Different models of an entity may represent different aspects of that entity (e.g. financial versus management structure of a company) or could represent contrasting alternative models of the same aspect of that entity.

As shown in Figure 1, each model is visually represented as an interactive graph, with entities displayed as nodes, and relationships between entities displayed as edges. Symmetric relationships are shown as undirected edges; all other relationships are shown as directed edges. Users may drag entities around, in order to visually group or separate them, and may zoom, translate, or rotate the graph to change the visual emphasis.
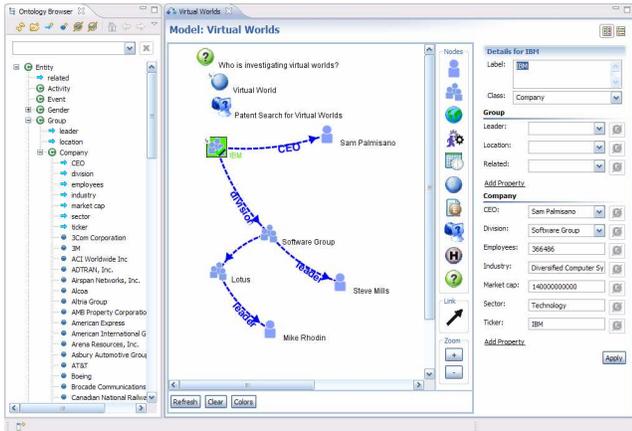

Figure 1: Screenshot showing a model that contains several entities. Details for the selected entity are displayed on the right; the shared ontology is displayed on the left.

Each asserted relationship is considered a claim, as are string and numeric property values asserted for an entity. Every claim has associated metadata, including the provenance of the claim, an analyst-asserted confidence in the claim, and any evidence that has been used to justify that claim. Some claim metadata is visually rendered in the corresponding graph edge. Figure 2 shows some of the possible edge styles; for example, negative claims are displayed in red, claims with high confidence are bold, and claims asserted without justification are shown as dashed lines, while those with justification are shown as solid lines.
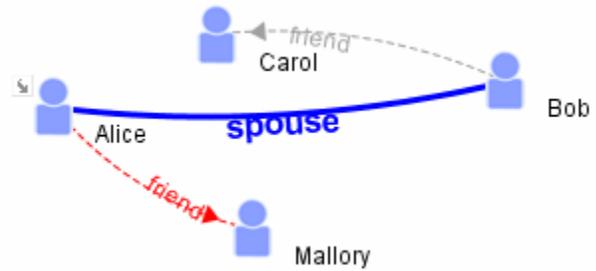

Figure 2: Model showing several edge styles. Bob is believed to be the spouse of Alice, and evidence is available for this symmetric relationship (solid, blue line). Carol might be a friend of Bob, but no confidence has been expressed (grey line). Mallory may not be a friend of Alice (somewhat negative confidence, hence the red line). The small arrow decorator indicates that Alice is a shared entity.

The entities and claims in a model are rendered as structured English text in a user interface component named the Progressive Summary. Each entity and claim generates a textual translation, and the translations update in real time as changes are made to the graphical model. Selection is synchronized between the nodes and edges in the graph and their translation. This provides an alternate representation of the graphical model and helps to enhance the understandability and accuracy of the model.

A palette of icons next to the model editor allows rapid addition of a new entity node via a drag-and-drop mechanism. Currently, icons are provided for a default set of useful classes, such as Person, Group, and Location classes. A planned enhancement is to make the palette configurable by the user, allowing quick access to commonly used classes. The palette also includes a link tool that allows the analyst to assert a relationship from one entity to another.

When an entity is first added to the model from the palette, it is considered a generic entity. The label of a generic node defaults to the name of the entity's class. If the class of a generic entity is changed, e.g. to a subclass of its current class, the label of the node is updated to match. Once the user has customized the label of the node, the entity is considered local and non-generic. If the analyst decides that a local entity is of general interest, it may be changed into a shared entity. Shared entities are available for use by all users of the CRAFT system, regardless of which investigations the user belongs to.

When a generic or local node is relabeled, we perform a case-insensitive search for shared entities with compatible membership and a matching label or alias. If one or more matches are found, the user is asked if the node should be replaced with a reference to a shared entity. If so, the local entity is removed, the shared entity is added in the same location, and claims contained in models of the shared entity are merged into the current model, preserving existing claims. In addition to helping the analyst specify a more complete model, this shared entity matching can alert

the analyst to cases in which disambiguation may be necessary – for example, if there are two distinct city entities with the same name. It also performs the function of alerting the analyst that others have an interest in this entity, providing an opportunity for conversation and collaboration.

When a node is selected, a form is displayed next to the model graph area that allows the user to edit details about the entity, e.g. its label, the classes it is a member of, and claim values for any properties whose domain is one of the entity's classes. For relationship (object) properties, the input control allows the user to select existing members of the property's range class from a dropdown list; an autocomplete feature for these values also helps the user quickly enter relationships. When an arc is selected, the user is given an opportunity to relabel the arc, and the appropriate property labels from the ontology are made available to ease the task. The user is also given the opportunity to examine the accumulated evidence supporting (or refuting) the corresponding claim, and to adjust the confidence associated with that claim.

Ontology Evolution

CRAFT users are empowered to extend the ontology to capture new concepts on demand as they conduct their investigations. During the course of an investigation, if an analyst wishes to represent a property or relationship between entities that is not currently supported by the ontology, or the analyst wishes to express that a particular entity belongs to a class that is not in the ontology, CRAFT prompts the user to confirm the addition of a property or class. As shown in Figure 3, the user is asked to provide extra information to situate the new class or property within the existing ontology, such as the property domain and range, or the new class's superclass. Reasonable defaults are provided based on the currently specified class of the entity.
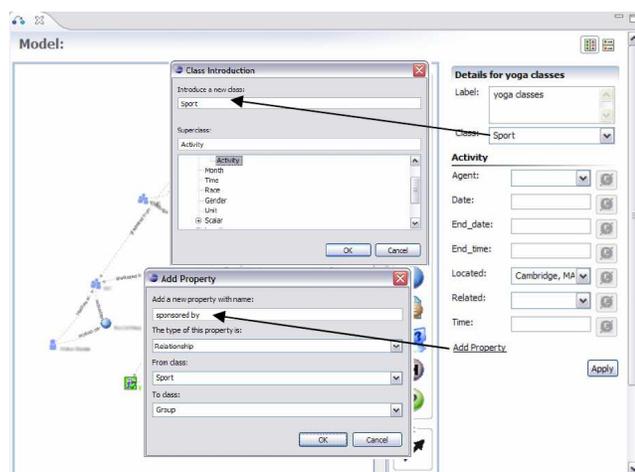

Figure 3: Introducing new classes and properties in CRAFT

Modifications to the ontology made by one user are immediately available for others to use, allowing the ontology to grow and evolve over time to reflect analyst's changing needs.

Because CRAFT is designed for investigational scenarios, we support additional node types that do not correspond to user-added classes in the ontology, such as Questions, Hypotheses, Evidence, and Inquiries. These types of nodes are treated differently from nodes corresponding to standard ontology classes. For example, we do not allow relationship edges to join these nodes to others.

Expressivity

CRAFT is built on top of Semantic Web technologies. Because it is designed for use by non-ontologists, the ontology features it exposes are limited - a proper subset of OWL Lite. (Bechhofer et al. 2004) The terminology surfaced in the user interface is chosen to minimize confusion for new users. For example, we refer to OWL individuals as entities because we found that users consistently interpreted "individual" to mean "person". Entities may be members of multiple classes, and classes may have multiple superclasses. Object properties may be marked functional, inverse-functional, symmetric, or transitive, but this functionality is only exposed through interfaces for advanced users. Currently only textual and numeric ranges are supported for datatype properties. CRAFT is not currently integrated with any automated reasoners, but the ontology can be imported or exported in RDF/XML format for use with other tools.

All semantic data and metadata is stored in Boca (Feigenbaum et al. 2007), a scalable open-source RDF named graph store. Boca provides support for multiple distributed users, replication and offline use, transactional updates with real-time notification to clients, and a SPARQL engine.

Ontology Navigation

As the size of the shared ontology grows, it becomes increasingly important to have methods for searching and browsing it.

One method for viewing the contents of the Ontology is a component named the Ontology Browser. This component displays the classes and inheritance hierarchy in a tree. (See Figure 1.) Properties are shown underneath their domain classes, and shared entities are also shown underneath the classes they are members of. The display of properties and entities in the Ontology Browser is optional. The Ontology Browser also allows the analyst to search for any resource (class, property, entity) in the continuously indexed ontology. Search results are provided as the user types, and include suggested synonyms from WordNet (Fellbaum 1998) for the entered search term. Buttons in this component's toolbar allow easy addition of classes, properties, and shared entities. Resources can be dragged from the Ontology Browser or

search results to a model. Dragging a class creates a new generic entity of that class, while dragging a shared entity adds a reference to that entity and imports claims from its primary model(s).

CRAFT also provides a component named the Resource Browser that provides detailed tables of information about any resource in the ontology, and that allows navigation from one resource to another using hyperlinks. The page displayed in the Resource Browser is synchronized with selections made in either the Ontology Browser or a model editor. In keeping with the webpage navigation metaphor, back and forward buttons are provided, and pages may be bookmarked for easy access. Other toolbar buttons allow modification of the viewed resource.
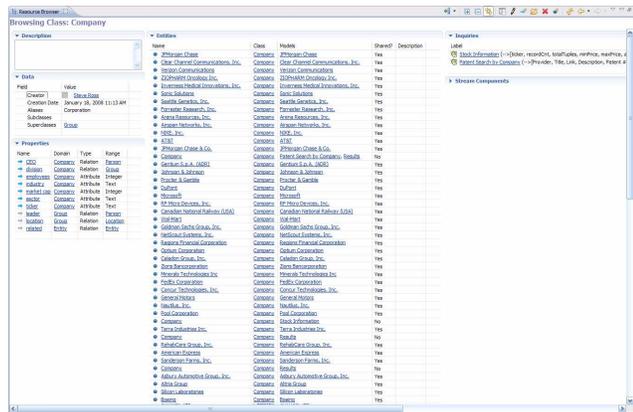
Figure 4: Resource Browser showing tables of information for the Company class.

Figure 4 shows the interface of an analyst viewing information about a "Company" class. There is a place to view freeform comments as well as structured metadata about the class's provenance. Aliases for the class and the placement of the class within the class hierarchy are also displayed and editable. Additional information includes tables of properties, entities, associated inquiries, etc. Icons indicate whether properties are inherited from superclasses or whether entities are direct members of the class or are members by virtue of belonging directly to a subclass. The user may choose to sort the table using the values of any column. As in a web browser, some hyperlinks will open a new window, such as a model editor, instead of changing the currently viewed resource.

We permit analysts to delete resources from either the Ontology Browser or Resource Browser. In most cases, the resource will only be marked deprecated and hidden from view. If CRAFT can determine that no references to the resource remain, the resource may be truly deleted to free computational resources such as disk space.

Inquiries and Analysis

An important part of CRAFT usage is the gathering of information and evidence by means of inquiries. When an analyst selects an entity node in a model, a selection of relevant inquiry templates is suggested. For example, a location node might suggest a weather inquiry template, while a company node might offer stock quotes or patent searches. These inquiry templates correspond to programs that can be run on the System S stream processing platform. (Amini et al. 2006) Unlike a web search or database query, CRAFT inquiries run continuously until explicitly stopped by the user. New results may potentially arrive hours, days, or even weeks after an inquiry is first submitted. Graphical indicators inform the analyst which investigations and models have new results, and the placement of an inquiry node within a model helps contextualize the results.

An analyst instantiates an inquiry template as a fully specified inquiry within a model by supplying the values for any required arguments. For example, a weather inquiry may require a zip code or city name, while an inquiry that searches a database of marriage licenses may have arguments for the social security numbers of each spouse. When the inquiry template is instantiated, we generate a dialog allowing the analyst to specify the values for the arguments. In order to help the analyst specify these values quickly and accurately, we allow an inquiry template to be associated with a semantic model of its arguments. This argument model, combined with the contents and selection in the target model, can be used to provide appropriate defaults for low-level argument values from claims in the knowledgebase.

CRAFT also supports another style of semantic inquiry, where the analyst creates a model using resources from the ontology to describe the desired result or pattern that should be matched.
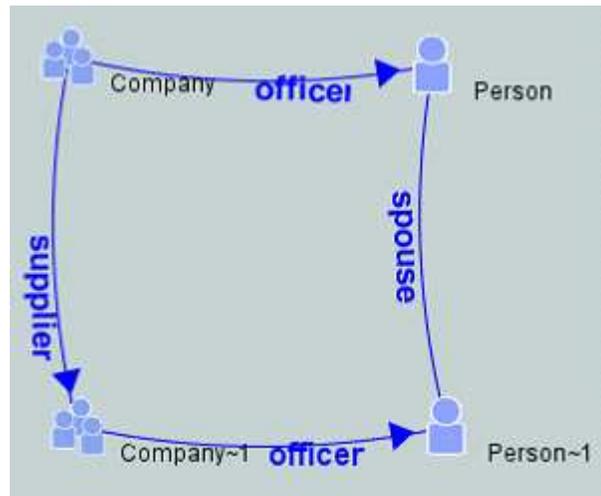
Figure 5: An example of a semantically expressed inquiry. This inquiry looks for potential conflicts of interest.

Figure 5 shows an example of a semantically expressed inquiry, where the analyst is searching for potential conflicts of interest in which an officer of one company is married to an officer of one of that company's suppliers.

In order to translate a semantically expressed inquiry like this to an executable program, a planner is used to compose a processing graph from semantically annotated streaming data sources and processing elements. (Liu, Ranganathan, and Riabov, 2007)

Analyzing Aggregate Knowledge

CRAFT is integrated with a variety of analysis and visualization tools. These may be applied to either the results of an inquiry or to a subset of the aggregate claims from the shared semantic repository. When applied to claims, any negative assertions are disregarded.

One of the analysis tools CRAFT supports is Exhibit (Huynh, Karger, and Miller 2007). CRAFT dynamically determines appropriate views and facets for a given set of data. For example, if one of the columns returned by an inquiry contains date information, a timeline view will be made available. (Figure 6) The facets displayed for filtering the dataset are also determined intelligently by examining the distribution of data values and excluding any facets that trivially categorize the dataset.
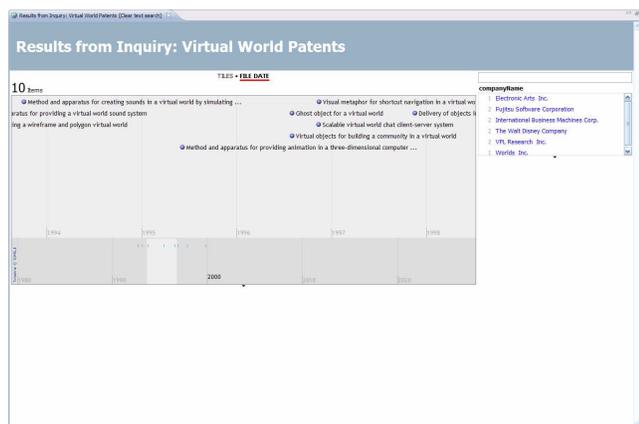
Figure 6: Exhibit used to explore the results of an inquiry. Here, ten patents about virtual worlds have been plotted on a timeline by their filing date. Facets allows filtering the patents by company name or by the presence of keywords.

CRAFT also supports analysis of inquiry results and claims with several of the interactive visualizations from Many Eyes (Viégas et al. 2007), such as pie charts, scatterplots, bubble charts, etc. Figure 7 shows an example of a Many Eyes visualization created from claims in the knowledgebase about the market capitalization of company entities. In the settings shown, each bubble's area is proportional to the cumulative market capitalization value for all companies in an industry. The bubbles are shaded according to industrial sector.
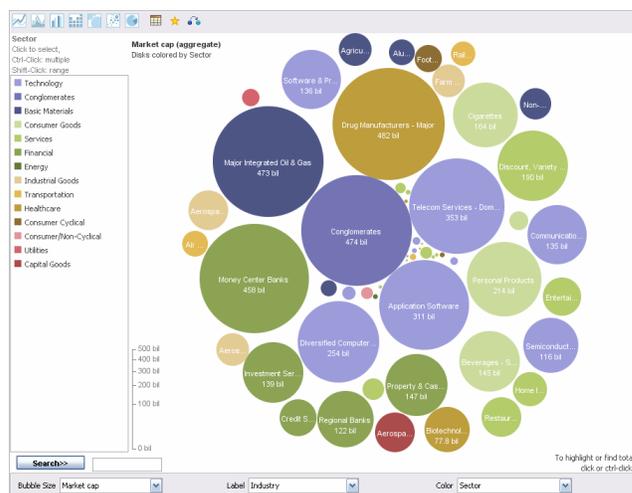
Figure 7: A Many Eyes Bubble Chart visualization created from claims in the knowledgebase

Entity selections within the visualization contain references to the models in which the claim values originate, allowing the user to drill-down from an overview visualization to the models containing the source data.

User Study and Evaluation

As described above, CRAFT supports a process of collaborative and implicit ontology evolution in which the ontology is modified on the fly to reflect new classes, properties, and relationships analysts wish to express. While this process has the potential to bridge the gap between knowledge engineering and usage, 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. We conducted a preliminary user study to investigate user behavior and ontology evolution in CRAFT. (Liu and Gruen 2008) Specifically, we explored the following 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 devised a paradigm in which multiple series of ontologies can be observed evolving in different trajectories from the same initial point. We analyzed and compared the different series of ontologies that evolved quantitatively using several published metrics. We also gained insights on user behavior around the ontology, and on the tool in general, through interviews and observations.

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, using an earlier version that lacked many of the summarization features described above. 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 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 in the same group created a model about another researcher R3 with the 1st generation ontology, resulting in the 2nd generation.

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

As shown in Table 1, each group created models for the same set of researchers, but in different orders. So we obtained three ontology series evolving on different paths from the same starting point.

We did not specifically bias the subjects to be careful in extending the ontology. The subjects were told to *"feel free to use anything existing and create anything necessary. The changes you make 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.

Summary of Results

1. Users were able to create meaningful models that were understandable by others.

In 30 minutes, the subjects created fairly complicated models, containing 13 to 26 nodes and 16 to 26 links, averaging 17.7 nodes and 20.9 links per model. A follow-up study showed that others were able to understand the models and extract information from them. Given the fact that this is the first time the subjects had used CRAFT, the result is encouraging. It shows that the subjects were able to effectively express their ideas using the tool.

2. Users were able to extend the ontology as needed, and the ontologies grew over time.

In the interview, all of the subjects said using the ontology was not constraining. As one subject stated: *"I can always create something new if it is not there"*. This suggests that integrating ontology extension seamlessly with end user tasks is a promising approach to addressing constraints of ontology enabled semantic applications.

3. There was a tendency to create "basic level" classes rather than abstract superclasses.

Our analysis of inheritance richness showed a tendency to create the basic level classes needed to represent specific items in an investigation, rather than more abstract, higher level superclasses. The ontologies became flatter as they evolved. It should be noted that in the prototype and tasks used in our study, there was little facility or need for summary views and visualizations by category. The value of defining higher-level superclasses was therefore less than it might otherwise have been, so there was less incentive for the users to do so. Users added meaning to the relationship arcs, using multiword relationships such as "used to work with", and adjusting them—as one subject reported—to "read right". This was no doubt influenced by the fact that in the prototype used in the testing, relationship arc descriptions were always displayed along the arcs, while only specific labels (e.g. "Bill"), and not class names (e.g. "Professor") were displayed for entity nodes. Users were using the facilities available to them to create a graphical representation that would communicate as much information on the surface as possible.

4. The three ontologies were different, but similarity increased over time.

Results of an analysis of Lexical Similarity (LS) and Term Overlap (TO) showed the three ontology series about the same domain were different from each other. However, as more subjects added more concepts, the similarity increased over the generations. An interesting phenomenon to note in Figure 8 is that the LS and TO for classes are very close, but not for relationships. Subjects had relatively high agreement on names of similar classes, but lower agreement for similar relationships. As discussed above, the names of relationships were more explicit and longer, so they are prone to more variation. Another reason for lower agreement relates to the direction that different subjects chose for basically the same relationship, e.g. "employer of" vs. "employed by".
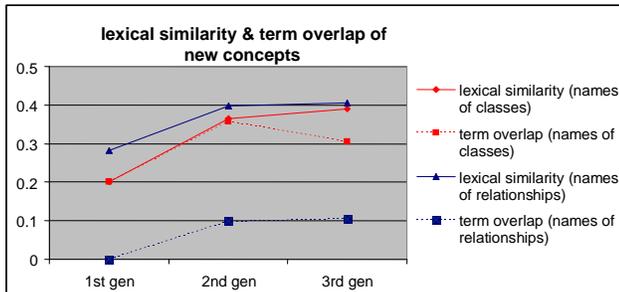
Figure 8: Similarity between the different ontology series at each generation.

5. Users were primarily influenced by their immediate investigational needs when modifying the ontology.

Not surprisingly, users were motivated more by their own needs in conducting their investigation than by considerations of the quality and general usefulness of the ontology that arose. We asked users: *"When adding new items, how much did each of these influence your choice?"*, and provided three options:

1. *Usefulness for other people using the ontology in the future*
2. *How well it fits in the specific model*
3. *How well it fits in with the other terms existing in the ontology*

Users overwhelmingly reported that the desire for it to fit into the specific model they were constructing (option 2) was their primary concern.

6. The graphical model was an effective tool for communicating information, and particularly valuable for going back to items seen earlier.

We also wished to evaluate the usefulness of the graphical modeling techniques used in CRAFT as an effective medium for communication. To do this, we compared the comprehensibility of graphical models with text reports by selecting three models of similar complexity in terms of number of nodes and number of links beyond those to the central node, and then created a textual report with the same information as in the models. Subjects were asked to "read" a graphical model and a text report of another model, and were then given a memory test (free recall after distractor task), and an information identification test, in which they were asked to answer specific questions and locate specific pieces of information as evidence. Results showed similar usefulness for information retention, with four subjects performing better with the graphical representation, and five better with the textual one (differences in performance were small). On the information identification task, however, eight of the nine subjects were faster with the graphical representation than with the textual one. There was one interesting case of a misunderstanding based on graphical configuration, in which a reader shown a graph created by another *assumed*

at first glance that nodes linked from one node were children, not siblings, as the layout looked like a family tree. Six of the nine subjects reported that they preferred the graph for information absorption, while all preferred the graphical model for information identification.

Summary

In this paper, we presented CRAFT, a tool for collaborative investigation, reasoning, and analysis. We described CRAFT's knowledge representation and modeling features, including their foundation on Semantic Web technologies such as RDF and OWL. We discussed some of the design tradeoffs we have made to balance expressivity and usability for analysts, our target user group, and how those analysts can find and reuse resources in a shared, evolving ontology. Two forms of semantically-based inquiries were described, as were two methods for analyzing and visualizing the results of those inquiries and amassed claims from the shared knowledgebase.

We also presented results from a preliminary study of collaborative ontology evolution using CRAFT. In a system such as ours in which users are empowered to modify an ontology on the fly as they conduct their work, specific details of the tasks they are doing and of the interfaces and tools with which they are doing their work have significant effects on the shape of the ontology that evolves. Based on these early but encouraging results, we intend to continue evaluation and development of CRAFT in order to provide continued value for our users.

Acknowledgments

References

Amini, L., Jain, N., Sehgal, A., Silber, J., and Verscheure, O. 2006. *Adaptive Control of Extreme-scale Stream Processing Systems*. Proceedings of the 26th IEEE International Conference on Distributed Computing Systems. icdcs p. 71.

Bechhofer, S., van Harmelen, F., Hendler, J., Horrocks, I., McGuinness, D. L., Patel-Schneider, P. F., and Stein, L. A. 2004. *OWL Web Ontology Language Reference*. W3C Recommendation, 10 February 2004. Available at http://www.w3.org/TR/owl-ref/.

Feigenbaum, L., Martin, S., Roy, M. N., Szekely, B., and Yung, W. C. 2007. *Boca: an open-source RDF store for building Semantic Web applications*. Briefings in Bioinformatics. 8: 195-200.

Fellbaum, C. (Ed.). 1998. *WordNet: An Electronic Lexical Database*. MIT Press.

Huynh, D. F., Karger, D. R., and Miller, R. C. 2007. Exhibit: Lightweight Structured Data Publishing. In *Proceedings of the 16th International Confernece on World Wide Web*. WWW'07. ACM, New York, NY.

Liu, J. and Gruen, D. 2008. Between Ontology and Folksonomy: A Study of Collaborative and Implicit Ontology Evolution. In *Proceedings of the 13th International Conference on Intelligent User Interfaces*. (IUI 2008). ACM, New York, NY.

Viégas, F., Wattenberg, M., van Ham, F., Kriss, J., and McKeon, M. 2007. Many Eyes: A Site for Visualization at Internet Scale. In *IEEE Transactions on Visualization and Computer Graphics*, vol. 13, no. 6, pp. 1121-1128.

Liu, Z., Ranganathan, A., and Riabov, A. 2007. *Use of OWL for describing Stream Processing Components to enable Automatic Composition*. Proceedings of the OWLED 2007 Workshop on OWL: Experiences and Directions. Innsbruck, Austria, June 6-7, 2007.