CRAFTing an Environment for Collaborative Reasoning

Susanne C. Hupfer, Steven I. Ross, Jamie C. Rasmussen, James E. Christensen, Stephen E. Levy,
Daniel M. Gruen, John F. Patterson

IBM T.J. Watson Research Center, One Rogers Street, Cambridge, MA 02142

firstname_lastname@us.ibm.com

ABSTRACT

We motivate the need for new environments for collaborative reasoning and describe the foundations of our approach, namely collaboration, semantics, and adaptability. We describe the CRAFT collaborative reasoning interface and infrastructure that we are developing to explore this approach.

ACM Classification: H.5.3 [Information Interfaces and Presentation]: Group and Organization Interfaces – Computer-supported cooperative work, Collaborative computing.

Keywords: Collective intelligence, intelligent interface, computer-supported cooperative work, collaborative reasoning environment, sensemaking, semantics, ontology.

INTRODUCTION

Many challenging problems ─ in diverse fields such as business intelligence, risk analysis, fraud detection, law enforcement, intelligence analysis, financial forecasting, epidemiology, and strategic planning ─ require sophisticated reasoning and are beyond the capacity of a single person to solve independently. These problems tend to be extreme examples of sensemaking activities ─ requiring the effort and knowledge of a variety of people with diverse expertise, collecting and sharing information and reasoning together over time to arrive at a common understanding of complex, changeable situations.

Sensemaking has been described as "a motivated, continuous effort to understand connections (which can be among people, places, and events) in order to anticipate their trajectories and act effectively" [7]. A number of systems have been implemented over the years that can be applied to sensemaking, even if not explicitly designed to support it. These include systems for information gathering and organizing, such as wikis, collaborative tagging systems, and collaborative information retrieval tools [4,13]; systems for information representation, idea generation, and brainstorming, such as hypermedia systems and concept mapping tools [1,6,12]; and systems for group argumentation and decision making [3,8,10,13].

Despite this array of existing tools, there is a class of sensemaking problems for which these tools, taken individually, are still inadequate. Other researchers have identified a number of factors that lead us to this claim [7]: Some problems lack well-defined start and end points, and may involve analysis of an uncertain situation that may evolve indefinitely. Some problems involve retrospectively examining past situations and events to arrive at an understanding, while others require anticipating what may occur in the future. Some questions and situations aren't amenable to a single definitive answer or decision, but require exploration of alternative explanations, which may need to be refined over time as supporting or refuting data are discovered. Some "wicked problems" are poorly understood, may have ambiguous or changing requirements, and may have contradictory solutions [2]. Moreover, the huge volume of information now accessible electronically (e.g. from the WWW, transaction systems, sensor networks) changes the way people approach complex issues and introduces its own issues of information overload and relevance. Indeed, many problems are too vast for any one person to take on alone – either there is too much information to sift through, or the expertise required is too far-ranging. We term these kinds of problems *collaborative reasoning* problems, and refer to frameworks for tackling them as *collaborative reasoning environments*.

In the next section, we present some foundations that we posit are needed for effective collaborative reasoning environments. We then introduce CRAFT (Collaborative Reasoning and Analysis Framework & Toolkit), a research project at IBM that we have designed and implemented to incorporate these capabilities. Finally, we summarize our experiences with CRAFT and discuss future directions.

COLLABORATIVE REASONING ENVIRONMENTS

Our goal was to develop an intelligent interface and underlying infrastructure to support people (working separately and together) in gathering and analyzing information and making decisions, while benefiting from the information-gathering and decision-making activities of others. Here we list elements we believe to be essential for such a system to be effective.

Collaboration

Some problems demand a collaborative approach. People cannot tackle securities fraud, global climate change, or homeland security alone. Moreover, researchers have demonstrated circumstances in which collaborative reasoning is superior to individual reasoning [11].

With collaborative reasoning problems, the objective of any investigation is merely a vision. As users gather data, the "goal" may evolve. Utilizing information gathered and assessed by other individuals aids in validating the objective and may open up other avenues of inquiry that were not obvious at the outset of the investigation. The path from data to knowledge, understanding, and solution is iterative, rather than a linear progression [2,7]. Information and its accuracy may change over time, which may prompt an investigator to revisit and reassess relevant hypotheses. Reasoning done by others may impact an investigator's own reasoning and prompt additional reassessments. Furthermore, knowledge of other investigations and their progress may lead one into new investigations.

Other important aspects of collaboration are presence awareness ─ knowing who else is currently working on the same topics or investigations ─ and expertise location ─ identifying relevant experts who can be consulted on a particular subject or investigation. Collaboration motivates the need to vet and qualify information, and to understand the social structures that parallel the information structures.

Semantics

Our point of view is that analysts need to create and share knowledge representations in order to effectively reason with others about information and situations. In our approach, investigators build semantic models that represent people, places, events and their relationships; they represent what they know about a situation, record questions, hypotheses, and evidence, and create inquiries for new information from databases and external sources ─ using concepts and instances from an evolving ontology.

Unlike most conventional approaches to sensemaking – such as typical wikis, collaborative tagging systems, and concept mapping tools – our approach is grounded in formal semantics. By tying into ontologies, the models that analysts create become a *lingua franca* for the exchange of information ─ both with other users and with the system ─ and begin to populate a shared knowledgebase. Since the semantic models are machine-interpretable, the system is able to assist users in ways it could not if analysts were just recording and organizing raw textual data. Use of semantics allows the system to help connect people to information and to one another. Suppose that while building a model an analyst creates an entity that is an instance of a class "Corporation" and termed "IBM". The system is able to detect whether an "IBM" corporation already exists in the ontology, and can ask whether the user wishes to use the existing entity. If the analyst does reuse the entity, all existing knowledge associated with that entity's model ─

e.g. that "Sam Palmisano" is the CEO and that the headquarters location is "Armonk" ─ is obtained for free. This information, collected by another analyst in some other context and possibly for a completely separate investigation, can now be leveraged and repurposed ─ a kind of serendipitous collaboration. Semantics thus helps us avoid duplication of effort. Semantics can also lead us to discover people with overlapping interests or a particular expertise; in the example above, the system can reveal who the expert on "IBM" is, and the two analysts may converse and/or collaborate more explicitly. Furthermore, aspects of the work the original analyst conducted may suggest lines of inquiry that the current analyst has not yet considered, and the latter may build upon the models and inquiries contributed by the former.

In theory, formal semantics could also allow machine reasoning to assist human reasoning. For example, the system could automate some aspects of data analysis, such as intelligent filtering of streams of information or discovery and composition of related models.

Adaptability

Because many of the collaborative reasoning problems that motivated this research are open-ended and ever-changing, we believe that collaborative reasoning environments need to support ongoing, long-term investigations. Unlike other sensemaking tools that support structured decision-making or argumentation ─ supporting users who focus on sharing knowledge, reaching a decision, and then moving on ─ collaborative reasoning tools need to be a much more integrated, adaptable part of an analyst's ongoing routine.

Collaborative reasoning environments need to support change and evolution, since they are intended for reasoning about a world that is always changing. In our approach, analysts create semantic models of problems and continue to revise them in response to new information from external sources and other analysts. Even long after an investigation is considered "cold," a new, relevant information item may enter the system or an existing item may be revised (e.g. a revised weapons assessment or new climate change data); applicable investigators should be made aware of it, and they may want to revisit the old investigation.

We are long past the days when all relevant information can be "poured into" the tool at the outset. A collaborative reasoning environment must be able to interoperate with information in the world-at-large, including structured and unstructured information, data and meta-data. Users should be permitted to collaboratively construct and issue *persistent contextual inquiries* from the relevant context within a model. Results should collect within that context as well. In this manner, results will accumulate where they are most relevant, automatically organizing the incoming data.

Unlike traditional centralized approaches to ontology development ─ in which ontologies are first built, then deployed, and later onerous to change ─ we advocate "on-
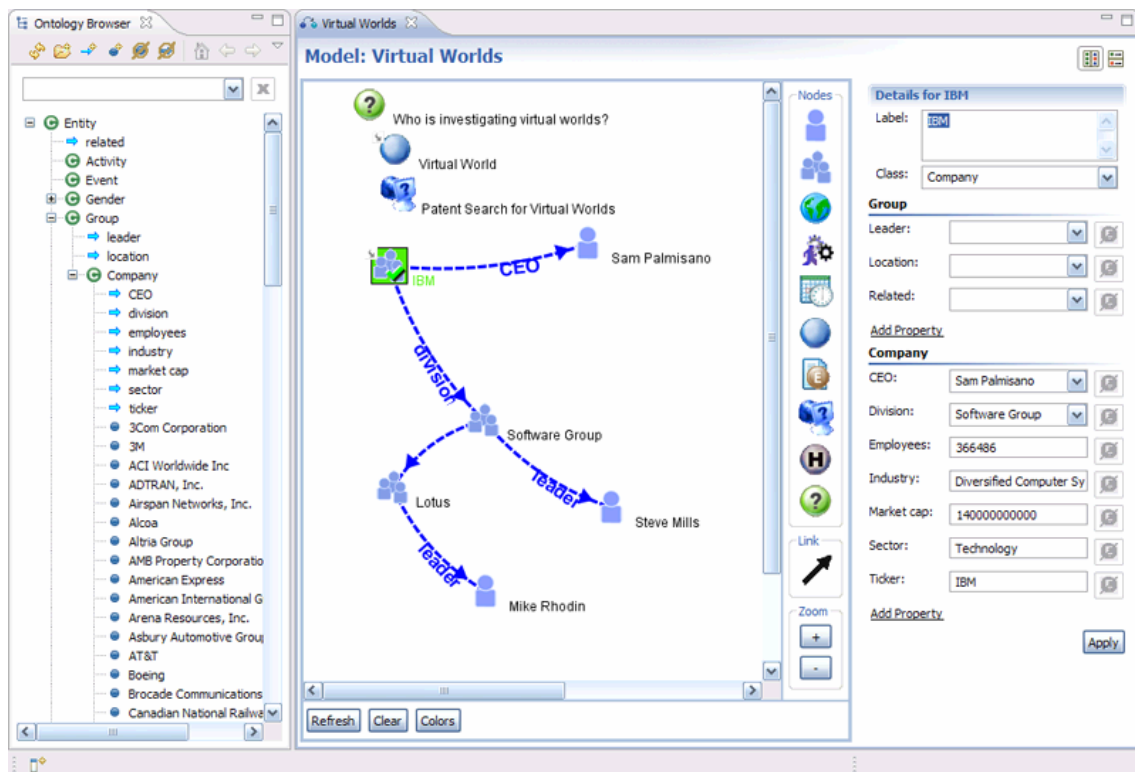
Figure 1. CRAFT user interface, showing Ontology Browser, Graphical Model Editor, and Details Pane.

demand collaborative ontology development." As users conduct investigations, they should be able to extend the ontology, capturing new classes and properties as needed. As models and ontologies evolve over time, changes made by each user should be made available immediately for others to use, resulting in community resources that grow and adapt organically to their users' needs.

Analysts are subject to change as well; they may focus on one particular investigation for a time, and then move on to other pressing matters. As we mentioned above, semantics can be valuable in helping direct users to relevant experts that have previously explored the same topics.

CRAFT: COLLABORATIVE REASONING & ANALYSIS FRAMEWORK AND TOOKIT
The Collaborative Reasoning and Analysis Framework and Toolkit (CRAFT) is a research prototype created at IBM to aid analysts as they gather, share, and reason about information; it is our first attempt to embody our vision of collaborative reasoning presented in the last section [5].

Analysts can belong to multiple ongoing investigations — each a kind of team space for members to work together to achieve a goal, answer a question, or research a situation. An investigation can contain one or more models, each of which contains references to entities (instances of ontology classes) and claims about those entities. Figure 1 shows the CRAFT interface. The center pane shows a graphical model editor that allows users to visually represent a model as an interactive graph, with entities shown as nodes and relationships shown as edges. Each entity may itself have

one or more associated models; e.g. an investment model might reference an IBM entity of class Corporation, and that entity may link to alternate models of IBM, e.g. a financial model and an organizational model. CRAFT supports presence awareness; analysts can see who is using the environment and what models they are working on.

Users can add new entity nodes to the graph via drag-and-drop from the model editor's icon palette, which provides common classes, e.g. Person, Group, Location, Activity, Event. They can also drag-and-drop classes listed in the Ontology Browser to the editor to instantiate entities. The model editor functions as a shared whiteboard; as users add, rearrange, label, and delete nodes and edges, others will see the effects. When nodes are selected, a Details Pane appears at the right; the form allows editing of entity details, e.g. label, classes it belongs to, and values for properties pertinent to those classes. When first added to the model, an entity is considered "local," but if it may be of general interest, the analyst can mark it "shared." Shared entities are available for use by all analysts. Whenever a user adds a new node, the system checks for shared entities that may be a match (e.g. the label matches) and offers use of those instead. This matchmaking has several benefits: It allows automatic reuse of existing claims contained in models of the shared entity, it alerts analysts to cases where similarly named entities need to be disambiguated, and it provides an opportunity for collaboration among analysts that are interested in the same entities.

The icon palette also provides a link tool that allows the analyst to assert a relationship — a claim — from one entity

to another. Every claim has associated metadata, e.g. its provenance, analyst-asserted confidence in it, and any associated evidence (supporting or refuting).

As they conduct investigations, analysts are able to extend the ontology to capture new concepts on demand. If one wishes to represent a relationship that is not currently supported by the ontology or to express that an entity belongs to a new class, CRAFT prompts the user to confirm addition of a property or class. Ontology changes are immediately available to others, and the ontology evolves over time to reflect analysts' changing needs.

CRAFT provides additional node types to support investigations: Question, Hypothesis, Inquiry, and Evidence. Analysts may instantiate a Hypothesis node and start building a model for a particular scenario, gathering evidence for or against it. Collecting information and evidence via "inquiries" is fundamental to CRAFT. When an analyst selects an entity, a selection of relevant inquiry templates is offered, e.g. a Location node might suggest a weather inquiry template, while a Corporation node might suggest stock quotes; inquiries correspond to programs that can be run on a stream processing platform or queries that can be set up to search databases or the web. Inquiries run until a user explicitly stops them. Visual cues notify analysts when investigations and models have new results, which may potentially arrive hours, days, or even weeks after an inquiry was first submitted. Analysts may create Evidence nodes populated by selected inquiry results, and then drag those nodes onto edges to support their assertions.

CONCLUSIONS AND FUTURE DIRECTIONS
While our work builds upon earlier research into sensemaking systems, the three pillars of our approach ─ *semantics*, *collaboration*, and *adaptability* ─ together make our viewpoint rather unique. Previous sensemaking systems have not made heavy use of semantics. Others have been designed for individual use, and have not placed an emphasis on collaboration. Yet others, while supporting collaboration, are designed for problems with distinct start and end points, and not geared toward continuous streams of incoming information or evolving structures [8,10,12].

We have implemented a fully functional CRAFT rich-client prototype in the form of Java extensions to Eclipse. We have conducted a preliminary user study to investigate user behavior and ontology evolution in CRAFT; in particular, the study focused on exploring the value of the *semantics* and *collaboration* aspects of our approach [9]. That study revealed that the graphical model is an effective tool for communicating information and valuable for revisiting items seen earlier, that users were able to create meaningful models that were understandable by others, and that users were able to extend the ontology as needed.

In order to continue evaluation of our approach, we are now moving toward a new web-based interface for CRAFT; we feel that this will make the tool more easily available to users beyond our own research group, and create a better infrastructure for further user testing. We hope in the future to be able to arm users with our tool and set them loose on their own challenging collaborative reasoning problems.

REFERENCES
1. Canas, A., Hill, G., Carff, R., Suri, N., Lott, J., Gomez, G., Eskridge, T., Arroyo, M. and Carvajal, R., CmapTools: A knowledge modeling and sharing environment. In *Proc. of the First International Conference on Concept Mapping,* (2005), 125-133.

2. Conklin, J. Wicked problems & social complexity, Chapter 1 of *Dialogue Mapping: Building Shared Understanding of Wicked Problems*, Wiley (2006), 1-20.

3. Conklin, J., and Begeman, M. gIBIS: A hypertext tool for exploratory policy discussion. *ACM Trans. on Office Information Systems 4,* 6 (1988), 303-331.

4. Golder, S. and Huberman, B. The structure of collaborative tagging systems. Technical report, Information Dynamics Lab, HP Labs, 2005.

5. Gruen, D., Rasmussen, J., Liu, J., Hupfer, S., and Ross, S. Collaborative reasoning and collaborative ontology development in CRAFT. *2008 AAAI Spring Symposium*, Tech. Report SS-08-07, 2008.

6. Halasz, F., Moran, T., and Trigg, R. Notecards in a nutshell, In *Proc. SIGCHI/GI Conf. on Human Factors in Computing Systems,* ACM (1987), 45-52.

7. Klein, G., Moon, B. and Hoffman, R.F. Making sense of sensemaking I: Alternative perspectives. *IEEE Intelligent Systems 21,* 4 (2006), 70-73.

8. Kraemer, K. and King, J. Computer-based systems for cooperative work and group decision making. *ACM Computing Surveys, 20,* 2 (1988), 115-146.

9. Liu, J. and Gruen, D. Between ontology and folksonomy: A study of collaborative and implicit ontology evolution. In *Proc. 13th International Conf. on Intelligent User Interfaces*, ACM (2008).

10. Lowrance, J., Harrison, I., and Rodriguez, A. Structured argumentation for analysis. In *Proc. of the 12th International Conf. on Systems Research, Informatics, and Cybernetics* (2000), 47-57.

11. Moshman, D., and Geil, M. Collaborative reasoning: Evidence for collective rationality. *Thinking and Reasoning, 4,* 3 (1998), 231-248.

12. Prante, T., Magerkurth, C., and Streitz, N., Developing CSCW tools for idea finding ─ Empirical results and implications for design. In *Proc. CSCW 2002*, ACM (2002), 106-115.

13. Romano, N., Nunamaker, J., Roussinov, D., and Chen, H. Collaborative information retrieval environment: integration of information retrieval with group support systems. In *Proc. 32nd Hawaii International Conference on System Sciences*, IEEE (1999).