The Giant: A Classroom Collaborator

Thomas R. Reichherzer, Alberto J. Cañas, Kenneth M. Ford, Patrick J. Hayes

Institute for Human & Machine Cognition
The University of West Florida
11000 University Parkway
Pensacola, FL 32514
treichhe@ai.uwf.edu

Abstract

The purpose of this paper is to report on the research involved in the development of a software agent—the Giant —that collaborates with students in their construction and sharing of knowledge. The Giant's working environment is an educational software tool used by students in and between classrooms to collaborate on a subject of study. Its task is to reason about the student's beliefs and ideas and automatically generate tentative conclusions. The Giant uses these conclusions to display its own, naïve understanding of the world to the student and asks him to decide upon its soundness. When interacting with the system, students assist the Giant in completing and correcting its propositions. This leads the students to review their personal, other students', and the Giant's claims, encouraging them to elaborate on their own knowledge.

Introduction

In Quorum, a partnership between The University of West Florida and IBM Latin America, a communication network was established to connect K-12 schools throughout many countries in Latin America. Among the objectives of the network was to produce an infrastructure that allowed students to collaborate in the construction and sharing of models of their beliefs [Cañas et al. 1996]. In order to reach this goal, Quorum provided not only an international communication network, but most importantly, a methodology for collaboration consisting of software tools, curriculum material, workshops along with pedagogical and technical support. One of the Quorum tools, the Concept Map Editor (CMap. Figure 1 shows a screenshot of the system in English. However, the system has been deployed in Spanish and Portuguese for the schools in Latin America.), focuses on the idea of constructing and sharing knowledge through concept maps. A concept map is a mediating knowledge representation that visually represents concepts and their relationships held by an individual. Concept maps are useful to tap into a person's cognitive structure and externalize concepts and propositions [Novak & Gowin 1984]. The idea of concept maps originated from assimilation theory [Ausubel et al. 1978] where meaningful learning, as opposed to rote learning, results from linking new information with relevant preexisting cognitive structures ('plant.cmp' Window in Figure 1 shows a simple concept map about plants).

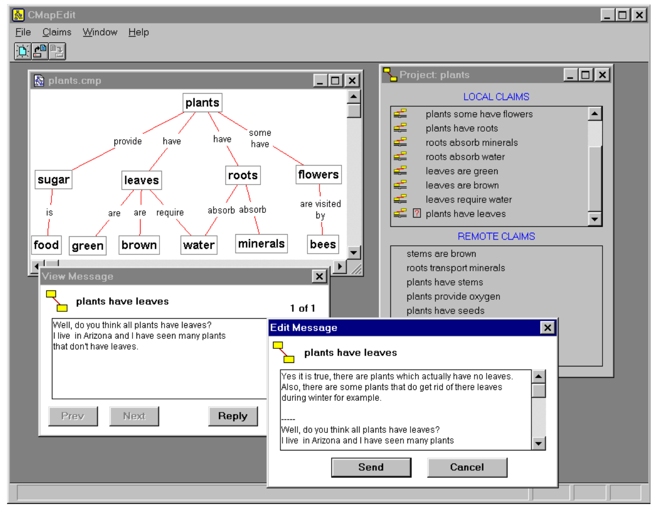

Fig 1. Knowledge construction and sharing using concept maps.

CMap's graphical interface offers students simple click and point operations to construct concept maps. As a student proceeds during the model construction, the system automatically identifies encoded propositions and stores them on the local workspace (see the "Local Claims" window in Figure 1). Optionally, students may identify propositions and extract them manually. By default, these propositions are local to the user and cannot be seen by other students. However, the system encourages collaboration among students during concept-map building by allowing them to share their propositions. By *publishing* (or *making public*) a proposition, the system adds it to a 'knowledge soup'—a collection of shared propositions stored in a server. At the same time, the system searches the knowledge soup for claims related to the student's local shared propositions and displays them (see the "Remote Claims" window in Figure 1) on the student's workspace (the system has built-in heuristics to decide upon relatedness). By publishing a proposition, the student "claims" that it is correct (hence the proposition is called a *claim*). Students can query other student's claims, if they disagree with it or don't understand it. Questioning a claim leads to discussion threads (see Figure 1, "View Message" and "Edit Message" windows) among those students who can view the particular claim. At no time does the system reveal to the users any information about ownership of claims. We believe that this aspect of anonymity will encourage users to step forward and share their ideas (see [Cañas et al. 1995] for a description of the knowledge sharing tools). A server that maintains the knowledge soup and distributes discussion threads and propositions can be located anywhere on a network. Through this system students from different classrooms, schools and from remote regions can collaborate during knowledge construction (Figure 2).

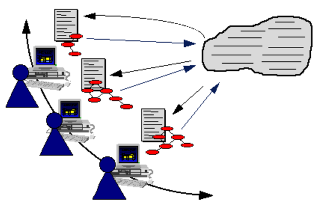

Fig. 2. Knowledge sharing across classrooms, schools and regions.

The Giant

The environment described is ideal for implementing a software agent that provides another source of collaboration to the students. The agent, embedded into CMap and running independently on each student's machine, generates tentative conclusions derived from both local propositions and shared propositions from the 'knowledge soup' by using a simple set of rules. It presents its own claims to the student and asks for a decision upon their soundness. By having to determine whether the agent's claims are correct, incorrect, or plain silly, students are placed into the position of a teacher to their agent, forcing them to review the agent's propositions, in addition to their personal propositions and the shared propositions from others. The agent uses the student's responses to improve its 'understanding' about the world. Additionally, students can question the agent to obtain an explanation for the agent's reasoning.

We refer to this agent as the Giant because in a sense it 'knows' a great deal but it completely lacks common-sense knowledge and has limited reasoning capabilities that sometimes induce a silly but amusing behavior. The system is not guaranteed to draw rational conclusions from the concept map (as a human being usually would) and in no way does it verify the student's propositions—such an intention is not pursued in our study. However, the Giant's propositions often act as a 'destabilizer' to the student, proposing conclusions that encourage the student to consider a different line of thought.

Behavior and appearance of the Giant

The Giant is equipped with an artificial 'personality' that presents it as a friendly, eager learner. It tries to capture the student's attention by presenting its understanding of the subject of study on which the student works, based on ideas and beliefs expressed in the student's concept map and shared in the classroom. The simplicity of its reasoning sometimes leads to less intelligent or conceptually wrong conclusions. This leads students to viewing the Giant as a harmless, non-threatening collaborator, motivating them to assist him in its effort to learn. We enhanced this impression on the Giant by providing the system with cartoon interfaces that reflect its current frame of mind (Figure 3 depicts some of the Giant's facial expressions).

Students may control the Giant's behavior by selecting from two settings in the system—the Giant's activity and curiosity. From these settings, students may choose their Giant to be active or lazy, curious or cautious.

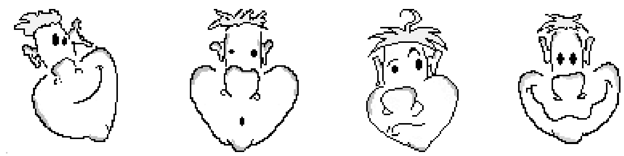

Fig 3. Facial expressions of the Giant.

An 'active' Giant is eager to produce as many propositions as possible whereas an 'inactive' Giant provides a small number of propositions that reflect its less active status. The Giant appears 'curious' about the student's work when it applies the entire set of rules to the propositions, yet does not insist on complete fulfillment of the premises of the rules. This behavior setting allows the Giant to derive conclusions that are not necessarily implied in the student's propositions. As a result the Giant might explore new facts about the subject of study. However, in this state there is a good chance that the Giant derives silly statements. The 'cautious' state tries to avoid this by executing rules only if they completely fulfill the premises.

The Giant has a general idea on how words affect concepts or relationships among concepts. These words, called keywords, play a significant role in the Giant's reasoning process. However, since the Giant has no access to common-sense knowledge or semantic understanding it cannot detect concepts or propositions in a student's map that might be appealing to question. In fact, the system selects randomly propositions from the student's workspace to perform its reasoning, causing it to appear moody or ignorant from time to time.

The Giant never intrudes upon the student's work. Rather, the student decides when to interact with it. Students may enter a dialog with the Giant by requesting an explanation for one of its propositions. At this point the system opens a dialog box to present the Giant's reasoning and to allow the student to either accept or refute the Giant's conclusion. Upon acceptance the Giant may request further information (see Figure 4) from the student to generate a new proposition. Students may use such information to extend their own concept map.

The Giant welcomes the student's advice and shows its appreciation when being taught. Such behavior is in contrast to the role of an authority that judges the student's work.

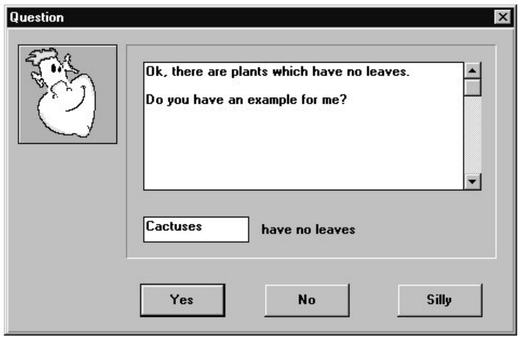

Fig. 4. The Giant requesting a new concept.

The Giant's reasoning

When generating claims, the Giant first parses the information obtained from both the student's local claims and the 'knowledge soup' using a grammar that tries to describe closely the syntax of propositions in concept maps. The result of this parse process is a representation of a proposition in machine-understandable form that the system applies to different categories of rules each capable of generating new propositions. This representation is used when the system matches or selects constituents of propositions in order to examine the premises of the rules or to construct new propositions. Heuristics, built into the system, indicate whether the production of a proposition is reasonable or not. Our heuristics are based on a list of words or word pairs, known to the system as key words, that signal the plausibility of a conclusion. Currently, this list is fixed, i.e. neither the Giant nor the user can modify the list. Later systems will explore automatic adaptation of the list using the user's feedback.

We designed the rules specifically to explore new concepts and new relationship among concepts with the student's help. Below we discuss the categories of rules that the Giant uses for its reasoning.

Transitivity

The first category of rules applies transitivity to propositions in the student's workspace. The rules try to match the right concept of a proposition with the left concept of another proposition. For example, given *plants have leaves* and *leaves are green*, the Giant concludes that *plants are green*. The rules check not only for matching concepts, but also for pairs of linking words that indicate reasonable conclusions, and the curiosity and activity status of the Giant. Only if the premise of a rule is fulfilled does the rule generate a new proposition.

Quantifier, Qualifiers, and Dependencies

The next category explores keywords that quantify or qualify the left concept or the action in the proposition. Such keywords are quantifying pronouns (for example: "some," "many," "a lot of"), cardinals, or adjectives associated with the left concept, adverbs associated with the verb in the link, or verbs that express causal dependencies among the concepts in the concept map.

A cautious Giant uses the rules only if a conclusion is plausible due to the existence of keywords in the proposition. For example, from *some stars are neutron stars* the Giant may conclude *that there are stars that are not neutron stars*, or *the ability to be neutron star requires something*, or *stars are not always neutron stars*. The simplicity of such conclusions shows the Giant's limited capabilities of sentence understanding and construction.

Classification / Extension

The last category of rules explores the structure of concept maps, searching for classifying concepts or new relationships among concepts. The rules trigger a response if a set of different concepts is linked with identical linking verbs to a common concept. The identical linking verbs and the common concept provide a hint that an enumeration of facts may have occurred. Figure 5 shows an example of a concept map to which the rule may be applied.

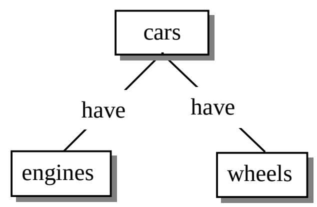

Fig. 5. Search for classifier concept.

This concept map enumerates parts of cars. One possible extension to the map is *engines and wheels generate motion* with "motion" as a new concept linked to "engines" and "wheels" through the linking word "generate." Another possible extension is "engines turn wheels" with "turn" as a new relationship between "engines" and "wheels." The classification/extension rules conclude that *there is a common identifier for engines and wheels* and *there exists a relationship between engines and wheels*.

Results

The tables below show the performance of the system on two small school projects with concept maps built in classroom settings by sixth-grade students. In particular, we were interested in the type of propositions and questions the Giant produces and how the students interact with the Giant. The 'knowledge soup' contained propositions derived from all the concept maps. The domains for the two sample projects are "the nature of plants" and "automobiles" respectively. Tables 1 and 2 present a subset of the student's and the Giant's propositions and consequent questions or conclusions by the Giant for each project.

In both tests, the student was actively collaborating with the Giant, trying to answer all of the Giant's questions. On average, the student refuted two out of ten questions as "silly" and, for three out of six questions, provided new concepts.

Student's propositions:	**Giant's propositions:**
some plants have leaves	there are plants that don't have leaves
Subsequent questions by the Giant:	
Do you know which plants do not have leaves? Do you know when plants don't have leaves?	

Student's propositions:	**Knowledge Soup claims:**
some plants have leaves	leaves make photosynthesis leaves release water
Giant's conclusion:	
some plants make photosynthesis some plants release water	

Table 1: Test results for project "the nature of plants."

Student's propositions:	**Giant's propositions:**
automobiles have wheels automobiles produce fume	there are automobiles that have not always wheels there are automobiles that do not produce fumes
Subsequent questions by the Giant:	
Do you know automobiles that have not always wheels? Do you know when automobiles do not produce fume?	

Student's propositions:	**Knowledge Soup claims:**
automobiles have engines	engines produce noise
Giant's conclusion:	
Automobiles produce noise?	

Table 2: Test results for project "automobiles."

Summary

We have developed an agent, called the Giant with working versions in Spanish, English, and Portuguese. The purpose of this agent is to take the role of a learner that seeks to understand the subject of study on which students in a classroom collaborate and shares with the students what it learns. Generating automatically tentative conclusions from the body of knowledge constructed by the students, it tries to motivate the student through its simple-minded personality to provide orientation. Our tests have shown that the Giant produces claims that mostly are interesting and smart but sometimes are boring and silly. Most important, however, is that the Giant's claims can suggest new ways of thinking which arise from other propositions. As a consequence, students need to analyze the Giant's propositions, in addition to their own and those of other students, leading them to refine and broaden their own knowledge. We anticipate that this process will result in richer and broader concept maps.

We have introduced a relatively small set of inference rules with encouraging results. Future efforts include to enlarge significantly the set of rules that perform the automatic reasoning, and to incorporate a consistency checker that would enable the Giant to test the consistency of the student's claims with the 'knowledge soup' as collaboration in the classroom progresses.

References

Allen, J. (1987). *Natural Language Processing*. New York: The Benjamin/Cummings Publishing Company.

Ausubel, D. P., Novak, J. D. Hanesian, H. (1978). *Educational Psychology: A Cognitive View* (2nd ed.). New York: Holt, Rinehart & Winston. Reprinted, 1986. New York: Warbel & Peck.

Cañas, A. J., Ford, K., Hayes, P. J., Brennan, J., Reichherzer, T. (1995). Knowledge Construction and Sharing in Quorum, *Proceedings of AI-ED 95*, Washington DC, pp. 218-225.

Cañas, A. J., Ford, K., Hayes, P. J., Brennan, J., Reichherzer, T., Suri, N. (1996). An Environment for the Construction and Sharing of Knowledge, *Proceedings of the Florida Artificial Intelligence Research Symposium (FLAIRS 96)*, Key West, pp. 238-242.

Novak, J. D. & Gowin, D. B. (1984). *Learning How to Learn*. New York: Cambridge University Press.