Automated Advice-giving Strategies for Scientific Inquiry

M. Paolucci[†], D. Suthers[‡], and A. Weiner[‡]

† *ISP, University of Pittsburgh, Pittsburgh, PA 15260*
‡ *LRDC, University of Pittsburgh, Pittsburgh, PA 15260*
E-Mail: paolucci/suthers/arlene+@pitt.edu

Abstract

We describe a prototype advisor for students using Belvedere, an environment for conducting discussions about scientific controversies. The advisor has two strategic components, syntactic and consistency-based. Syntactic strategies are based on structural and categorical patterns in argument representations constructed by the students, and suggest ways in which students can continue their inquiry. Consistency-based strategies check student-made links between pairs of statements against the pairwise relations specified between corresponding units in a knowledge base constructed by a teacher or expert, and identify information that may challenge or corroborate relationships proposed by the students.

1. Introduction

Consistent with research that shows that students learn better when they actively pursue understanding rather than passively receiving knowledge [3, 8, 15] the classroom teacher is now being urged to become a "guide on the side" rather than "the sage on the stage." New roles have been recommended for ITS that parallel the teacher's new role in "decentered" classrooms [2, 5, 9]. Our work is addressed towards ITS that augment the learning processes of students engaged in collaborative critical inquiry [7, 10].

With others in our group,[1] we have developed a computer environment that supports students in conducting discussions of problems of scientific knowledge. In this environment, called "Belvedere" [13, 12], students can explore problematic situations by searching on-line information and by producing and manipulating graphical representations of arguments called "inquiry diagrams."

We have now prototyped an automated advisor that gives advice on demand concerning ways in which an argument in this environment can be extended or revised. Rather than supplying oracular advice whenever the student missteps, the advisor is on-demand, avoiding inappropriate intrusion into student discussion that may be taking place external to the computer environment. Advice is phrased as suggestions and questions because we cannot presume that an automated advisor has sufficient information to be imperative, and we want students to think about the advice, not just execute it.

In this paper we discuss two methods of advice generation that we have implemented. Syntactic advice strategies make suggestions based solely on the syntactic structure of students' inquiry diagrams. Consistency-based advice strategies use a simple knowledge base of consistency relations between information units to identify information that may

---

[1] Violetta Cavalli-Sforza, John Connelly, Alan Lesgold, and Arthur Nunes; recently joined by Kim Harrigal, Dan Jones, Eva Toth, and Joe Toth.

challenge or corroborate relationships postulated by the students. Before describing these advice giving methods, we first briefly describe the activity that the students are engaged in.

2. The Scientific Inquiry Tasks

The Belvedere interface affords students the means to readily produce symbolic diagrams, by assembling graphical shapes and connecting them with a variety of links, typing in text, and searching an index of online documents to cut and paste text from them into the diagrams. The students select the shapes and links they want from a column of icons labeled with the names of argument components (see figure 1). In order to focus the students' attention on the argument relations, some of the work of drawing and layout is automated, such as sizing shapes to fit text and positioning of links.

In a typical session in lab studies, pairs of students worked side by side on individual monitors, close enough to see and point to each other's displays. Working together in a shared drawing space, they were asked to find information to resolve an actual scientific problem, and to use Belvedere graphics to express the work that they did together. Problems that we have used hinge on conflicts between scientific theories or between a theory and discrepant data. For example, in the "iguanas problem," the students begin with a text explaining an anomaly for evolutionary theory: the two species of iguanas native to the Galapagos islands would appear to have separated more than 12 million years ago, but the Galapagos islands are dated at only 2-3 million years of age, conflicting with the theory that the species diverged on the islands [1]. In the studies, the students were given a 3-node graphical representation of the problem, the tools for extending the representation, and access to brief on-line texts relevant to the problem (as well as some that were not relevant). They were asked to work together to try to resolve the anomaly. Results of our formative evaluations are described in [13, 12].

3. Pedagogical Constraints on Advice

Our design of the advisors to be discussed were guided in part by the following constraints.

Maintain the student-initiated character of Belvedere's environment. Belvedere encourages reflection by allowing students to see their argumentation as an object. They can point to different parts of it and focus on areas that need attention. They can engage in a process of construction and revision, reciprocally explaining and confronting each other. An advisor should not intervene prematurely in their thinking process. It should be discreet, offering advice on request. Students should feel free to discard an advisor's suggestions when they believe them to be irrelevant or inappropriate.

Address certain parts of the task that are critical to the desired cognitive skill. Research on "confirmation bias" and hypothesis driven search suggests that students are likely to be concerned with the process of constructing an argument for a favored theory they

Figure 1: Diagram by Students on HIV/AIDS Issue.

are supporting, sometimes overlooking or discounting discrepant data [6, 4]. Also, they may not consider alternate explanations of the data they are using. An advisor should address these problems. For example, it should offer information that the student may not have sought, including information that is discrepant with the student's theory.

Be applicable to problems constructed by outside experts and teachers. The advisor should be able to give useful advice based on a simple knowledge base that an expert or a teacher might construct. So far Belvedere has been used to construct arguments in domains as different as theory of evolution, contrasting theories of mountain formation, cause of the Cretaceous extinctions, whether HIV causes AIDS, and theories in social psychology. It is not feasible to develop for each a representation of the knowledge needed to deal with the argumentation students potentially could engage in. We are instead interested in a general approach, applicable to all the cases, in which the knowledge base can be constructed by a teacher.

4. Syntactic Advice Strategies

The first approach we implemented gives advice in response to situations that can be defined on a purely syntactic basis, using only the structural and categorical features of the students' argument graphs. (The students' text is not interpreted.) Types of advice are defined in terms of patterns to be matched to the diagram, and textual advice to be given if there is a match. Example advice patterns are given in Table 1.

```
(def-advice circular-support
    :advice "This looks like a circular argument. Is there a statement in this group of
             statements that doesn't depend on accepting the rest of them?"
    :arguments (?x)
    :test (:and (statement ?x) (support-or-explain* ?x ?x))
    :type :incoherence)
(def-advice support-competitor
    :advice "Could the empirical data that supports one theory also support the other?"
    :arguments (?t1 ?t2 ?d)
    :test (:and (theory ?t1) (theory ?t2) (empirical ?d) (:not (:same-as ?t1 ?t2))
           (supports* ?d ?t1) (:not (supports* ?d ?t2)))
    :type :open-world)
```

Table 1: Examples of Syntactic Advice Patterns

The advice applicable to a given inquiry diagram is often more than a student can be expected to absorb and respond to at one time. When more than one instance of advice is applicable, a preference-based quicksort algorithm is used, following a mechanism used by Suthers [11] for selecting between alternate explanations. Advice instances are sorted in priority order, and the highest priority advice is given. Objects that bind to variables in the patterns are highlighted in yellow when the advice is given, so the user can easily identify what the advice is about. If further advice is requested before the diagram changes, subsequent advice instances on the sorted list are used without reanalysis. We are investigating preferences that take into account factors such as prior advice that has been given, how that advice has been responded to, how recently the object of advice was constructed, and various categorical attributes of the applicable advice.

We believe that the most important kind of advice is that which stimulates and scaffolds constructive activity on the part of the students. To give this kind of "open world" advice, our first step was to identify partial argument patterns in the inquiry diagram the students had constructed so far and indicate how the student could complete these patterns. For example, the advisor might find theoretical claims that have no empirical support and suggest that support be sought, or it might find competing theories where one theory is supported by some empirical observation and ask if the same observation can support the other theory (support-competitor in Table 1).

The syntactic advisor also responds to illegal and incoherent constructions. "Illegal" constructions are those that use elements of the diagrammatic language in a manner inconsistent with their intended semantics. For example, a "support" link should not be used between data. "Incoherent" constructions are those in which the elements are each used legally, but in combinations that are semantically problematic. Examples include a loop of "support" links (circular-support in Table 1), or a datum that both supports and undermines the same claim.

Consider the inquiry diagram shown in Figure 1, constructed by two students during formative evaluations at a public school site. The following advice is given (one by one and in this order) by our syntactic advisor when applied to that diagram:

1. Highlights the unconnected hypothesis (lower right); asks whether the students can say how this part relates to the rest of the diagram.

2. Highlights the same hypothesis and asks whether the students can find a way to support it, or show that it predicts or explains an observed phenomenon.

3. Highlights the two theories and one of the "Data" boxes that supports one of them and asks whether it is possible that the same data supports the other theory. ("Data" is a kind of "Empirical Observation." This is an instance of support-competitor.)

4. Highlights the "explains" link (right of center), points out that ideas such as theories and hypotheses explain empirical observations, not the other way around, and tells the student an easy way to reverse the link.

In addition to those mentioned above, other kinds of advice that we have implemented include suggesting that a theory or hypothesis be formulated when none is present in the inquiry diagram; asking whether there is another theory that provides an alternate explanation for the empirical data when only one theory or hypothesis is involved in the inquiry diagram; and asking whether data can be found to discriminate between two theories that have identical support.

5. Consistency-Based Advice Strategies

Ideally, we would like to have an advisor that understands the students' text as well as the domain under discussion, and provides fully knowledge-based advice. This is not currently possible due to the difficulty of constructing domain knowledge bases and of understanding students' texts. Instead, we have adopted the strategy of investigating how much useful advice we can get out of a minimal semantic annotation before we move on to more complex approaches. In this manner we hope to better understand the cost/benefit tradeoff between knowledge engineering and added functionality.

The consistency-based advisor is our first step in this direction. It is intended to offer specific information that the student may not discover on her own. It makes two assumptions: students construct their inquiry diagrams from existing units of text, and these units are annotated with relationships recording whether they are consistent or inconsistent with each other, based on expert judgment. The advisor searches the latter "consistency graph" to find paths between units that students have used in their inquiry diagrams, and selects other units found along those paths which are brought to the students' attention. Our claim is that this enables us to point out information that is relevant at a given point in the inquiry process without needing to pay the cost of a more complete semantic model of that information.

5.1 The Consistency Graph and Algorithm

The algorithm for the consistency-based advisor is based on a comparison of the student's inquiry diagram with information derived from a teacher's or expert's inquiry diagram. For the purposes of advice-giving, Belvedere's various relations between argument components are classified simply as relations of inconsistency or consistency and are
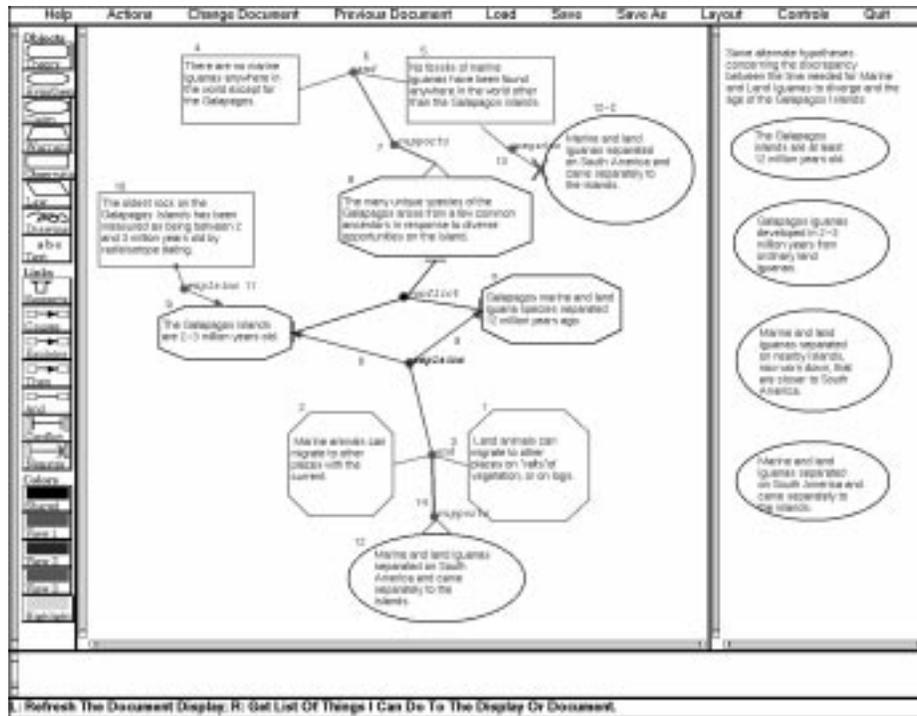
Figure 2: Diagram by "Emin" and "Mo," Galapagos Iguana Anomaly, early version of Belvedere

assumed to be symmetrical. Thus a "consistency" link means that the information in the connected nodes is at least compatible, and preferably one can be offered as evidence for the other. If a theory explains some datum, then the theory and the datum are "consistent." If a datum conflicts with a theory, then they are "inconsistent." In the diagram displayed in figure 2, the links "support," "causes," "explains" and "then" are interpreted as consistency relations. Links like "conflict" and "negates" are interpreted as inconsistency relations. The link "and" is a concept-forming link: it defines an implicit node that is consistent with the conjuncts, and inconsistent with nodes that are inconsistent with either conjunct. More precisely, the consistency relation of and-links is based on the following rules:

1. $(A \wedge B)$ is consistent with each of A and B individually.

2. If C is inconsistent with A or inconsistent with B, then C is inconsistent with $(A \wedge B)$.

The converse of rule 2 does not hold: $(A \wedge B)$ can be inconsistent with C while both A and B are individually consistent with C.

When the teacher or domain expert defines the task and the information needed, she can draw an inquiry diagram in Belvedere. The diagram is easily transformed into a

set of consistency relations between pairs of texts, following the rules described above. These relations become the "expert model" or knowledge base for the consistency-based advisor.

During a student session, students can drag texts authored by the expert into their diagrams, and express argumentation relationships between them. The links in the student diagram are interpreted as consistency relationships in the manner described above. The advisor can then compare consistency relations defined by the students with consistency relations defined by the expert, and look for inconsistencies and other possible advice. The comparison is based on a graph-search algorithm, which has been implemented and tested as reported below. It searches in the *expert's* consistency graph for paths between nodes that have been related in the *students'* graph, constrained by the following rules:

1. Only the shortest path between two nodes is considered.

2. A "positive path" crosses only consistency links.

3. A "negative path" ends with an inconsistency link.

4. A path can cross only one inconsistency link. An inconsistency link ends a path.

5. Conflicts between positive and negative paths are resolved in favor of the negative path.

The advisor can then select an item on the path found in the expert's graph that is not present in the students' graph, and present this to the student for consideration. If the path is of a different polarity than the students' link, the information presented could possibly contradict the relationship claimed by the students; if the polarity is the same the information would presumably support the relationship claimed by the students. (See figure 3 for an example.)

The five rules presented above define a non-monotonic logic similar to Thomason's skeptical reasoning [14]. Rule 1 is used to control the search and control the length of meaningless paths. The "consistency relation" is weaker than logical implication. It can easily be the case that inconsistent statements are at the ends of a long chain of consistency links. Although limiting the search to the shortest path does not solve this problem, it greatly reduces the effect of long and meaningless paths.

Rules 2 to 4 are used to maintain the consistency of the path. Once a negative link is crossed, it is plausible that the two nodes at the end of the path are inconsistent. We cannot extend the path any farther, because whatever conclusion is drawn from there is quite arbitrary. For example, suppose that A is inconsistent with B, and B is inconsistent with C. We can't conclude that A is inconsistent with C: they could be consistent components of an argument against B. On the other hand, we can't assume that A is consistent with C: they could be arguing against B based on incompatible assumptions. Thus, Rule 4 forces the search to stop when a conflict is reached. Rule 5 is introduced to address the confirmation bias.
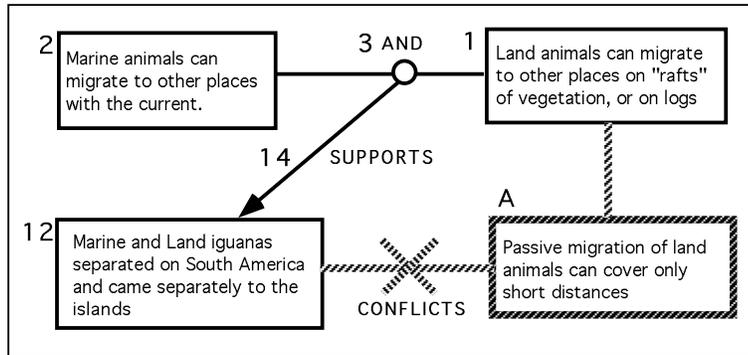
2 Marine animals can migrate to other places with the current.

3 AND

1 Land animals can migrate to other places on "rafts" of vegetation, or on logs

14 SUPPORTS

A Passive migration of land animals can cover only short distances

12 Marine and Land iguanas separated on South America and came separately to the islands

CONFLICTS

Figure 3: Example consistency path identifying information to be used in advice

5.2 Preliminary Evaluation

We have conducted two preliminary experiments with the consistency-based advisor. In the first experiment we were interested in testing consistency relations that we expected to be difficult or that required some inferential power. We used a subset of the "iguana problem" knowledge base used in some of the studies with the students comprised 19 nodes, 14 consistent and inconsistent relations, and 2 and-links. The three authors made judgments of consistency between pairs of statements corresponding to the nodes. Then we compared our judgments with the advisor's judgments. In all the relations about which all three authors agreed, the advisor made the same judgment. The only disagreements were on relations about which the authors disagreed. These cases were all characterized by the lack of a connecting path between the tested nodes. Either the search process was blocked by an inconsistency link, or a critical link was missing in an intermediate step of the process.

In the second experiment, we were concerned with the advice that would be given in a real interaction with students. We constructed a consistency graph of 90 statements and 73 relations from the materials used in one of the sessions with students, and performed consistency analysis on each link from two student sessions (see diagram in figure 2 for one of the sessions). The performance was similar to the previous experiment. We always agreed with the system's judgement, and the intermediate steps were sequences of coherent proofs. On most of the links the advisor agreed with the students (these were among our best students). In one case only, link 14, the advisor gave a different judgement: see figure 3. The path the advisor constructed starts at node 3, crosses node 1 and node A (not displayed in the students' graph), and ends at node 12. The advisor recognizes that this path (shaded) crosses an inconsistency link, so conflicts with the students' link 14. If the students would ask the advisor for a critique of their argument, the advisor would highlight link 14 and it would display Node A (the only information on the path that they have not seen), confronting them with the conditions for land animals' migration which they overlooked.

6. Conclusions

The syntax-based advisor can make suggestions to stimulate students' thinking with no knowledge engineering required on the part of the teacher or domain expert. However, the advice is very general and does not adequately address the confirmation bias. The consistency-based advisor can provide students with assistance in identifying relevant information which students may have not considered (perhaps due to the confirmation bias), and which may challenge their thinking. This advice cannot be provided by the syntax-based advisor, because the advice depends on knowledge of certain semantic relationships between the textual units involved. The level of "understanding" of the texts on the part of the system required is extremely minimal: this is an advantage, as it reduces the knowledge engineering demands on educators preparing materials for students. Clearly, a minimal semantic approach have limitations. For example it cannot infer the goals of the student, in particular which theory she is trying to build or support. The advisor described in this paper cannot help the student in the construction of an argumentation or find a counter argument that attacks her theory or engage the student in a scientific discussion. However, by investigating the utility of advice obtained from these minimal semantic annotations we hope to gain interesting insights that will help us to move toward more complex approaches, also we can better understand the cost-benefit tradeoff between knowledge engineering and added functionality.

References

[1] M. W. Browne. Galapagos mystery solved: Fauna evolved on vanished isles. *Science Times, the NY Times*, January 21 1992.

[2] T.W. Chan and A.B. Baskin. Studying with the prince: The computer as a learning companion. In *Proceedings of the International Conference on Intelligent Tutoring Systems*, pages 194–200, 1988.

[3] M.T.H. Chi, M. Bassok, M. Lewis, P. Reimann, and R. Glaser. Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science*, 13:145–182, 1989.

[4] C. Chinn and W. Brewer. Factors that influence how people respond to anomalous data. In *Proc. 15th Annual Conf. of the Cognitive Science Society*, pages 318–323, Hillsdale, NJ: Lawrence Erlbaum, 1993.

[5] W. J. Clancey. Guidon-manage revised: A socio-technical systems approach. In *Intelligent Tutoring Systems*, pages 21–36, 1992.

[6] J. Klayman and Y.-W. Ha. Confirmation, disconfirmation, and information in hypothesis testing. *Psychological Review*, 94:211–228, 1987.

[7] D. K. O'Neill and L. M. Gomez. The collaboratory notebook: A distributed knowledge-building environment for project-enhanced learning. In *Proceedings of Ed-Media '94*, pages 416–423, Vancouver, BC, June 1994. Charlottesville, VA: Association for the Advancement of Computing in Education.

[8] L. Resnick and M.T.H. Chi. Cognitive psychology and science learning. In M. Druger, editor, *Science for the Fun of it: A Guide to Informal Science Education*, pages 24–31. National Science Teachers Association, 1988.

[9] J. Roschell. Designing for cognitive communication: Epistemic fidelity or mediating collaborative inquiry? *The Arachnet Electronic Journal on Virtual Culture*, 2(2), 1994.

[10] R. E. Slavin. *Cooperative Learning: Theory, Research, and Practice*. Prentice-Hall, Englewood Cliffs, NJ, 1990.

[11] D. Suthers. Preferences for model selection in explanation. In *Proc. 13th International Joint Conference on Artificial Intelligence (IJCAI-93)*, pages 1208–1213, Chambery, France, August 1993.

[12] D. Suthers and A. Weiner. Groupware for developing critical discussion skills. In *CSCL '95, Computer Supported Collaborative Learning*, pages 341–348, Bloomington, Indiana, October 1995.

[13] D. Suthers, A. Weiner, J. Connelly, and M. Paolucci. Belvedere: Engaging students in critical discussion of science and public policy issues. In *AI-Ed 95, the 7th World Conference on Artificial Intelligence in Education*, pages 266–273, Washington, D.C., August 1995.

[14] R. H. Thomason. Netl and subsequent path-based inheritance theories. *Computers Math. Applic.*, 23(2-5):179–204, 1992.

[15] Noreen Webb. Peer interaction and learning in small groups. *International Journal of Education Research*, 13:21–40, 1989.

Acknowledgments