

Inquiry-Based Teaching and Learning of Science with Cognitive Assistants

Mihai Boicu, Gheorghe Tecuci, Dorin Marcu, James Trefil, Nancy Holincheck

Learning Agents Center and Advanced Studies in Teaching and Learning, George Mason University, Fairfax, VA 22030
{mboicu, tecuci, dmarcu, jtrefil, nholinch}@gmu.edu

Abstract

Inquiry-based teaching and learning is recognized as being very effective, but very difficult to use in practice. This paper introduces a computational framework of scientific inquiry as discovery of evidence, hypotheses, and arguments, which can be implemented in a cognitive assistant. With the help of two case studies, it shows how such a cognitive assistant that incorporates knowledge about the scientific inquiry process and about the properties, uses, discovery, and marshaling of evidence, can support inquiry-based teaching and learning.

Introduction

In this paper we address the problem of developing cognitive assistants for inquiry-based education. We first introduce teaching and learning of science through inquiry, and discuss how cognitive assistants have the potential of making this very promising educational approach both easier to implement in the classroom, and more effective. We then introduce a computational framework of scientific inquiry as discovery of evidence, hypotheses, and arguments, which can be implemented in a cognitive assistant. This framework is based on our previous work on developing cognitive assistants for intelligence analysis. We are building directly on the latest of these cognitive assistants, called Cogent (Tecuci et al., 2015), to design and develop cognitive assistants for inquiry-based teaching and learning of science.

To show the applicability and the generality of our approach, the rest of the paper presents two case studies of using cognitive assistants to support inquiry-based teaching and learning. The first case study is an adaptation of the classical example of using inquiry in the classroom from *Inquiry and the National Science Education Standards* (NRC, 2000). The second case study is based on our current work on developing *sInvestigator*, a cognitive assistant for helping undergraduate students develop critical thinking skills in addressing scientific problems.

We conclude the paper with the main directions of our future work.

Inquiry-based Science Teaching and Learning

Significant progress has been made in science education with the development of the National Science Education Standards (NRC, 1996). These standards call for inquiry-based teaching and learning which “refers to the diverse ways in which scientists study the natural world and propose explanations based on the evidence derived from their work.” Students practice inquiry as they “describe objects and events, ask questions, construct explanations, test those explanations against current scientific knowledge, and communicate their ideas to others. They identify their assumptions, use critical and logical thinking, and consider alternative explanations” (NRC, 1996, p. 2).

Researchers have demonstrated that academic achievement is improved by the use of inquiry instruction in K-12 levels (Bransford and Donovan, 2004; Minner et al., 2010). Inquiry instruction has also been examined at the college level and found to be more effective than traditional science instruction for the development of thinking and problem solving (Oliver-Hoyo et al., 2004). University science faculty value inquiry, but identify time, class size, student motivation, and student ability as obstacles to implementing inquiry-based instruction (Brown et al., 2006).

A significant result in the theory of inquiry-based learning is Process-Oriented Guided-Inquiry Learning (POGIL, 2016), a student-centered, group-learning instructional strategy and philosophy. POGIL provides a general framework for developing activities implementing guided inquiry in the classroom, and there are now many POGIL inquiry-based learning activities in a wide variety of disciplines. However, while POGIL and other class-activity based inquiry approaches offer an alternative to lectures-style instruction, they depend on intensive training of instructors to develop and implement inquiry-based activities in their classrooms.

How Can Cognitive Assistants Help?

We will briefly discuss some of the challenges of using

inquiry-based teaching and learning and how cognitive assistants may alleviate them.

Designing inquiry-based lessons. Even following a methodological framework (such as POGIL), the effort and time required to design such lessons are greater than those required for content-based lessons. Indeed, the teacher would not only need to introduce the theoretical topics but also to guide their discovery through inquiry-based activities. Even when inquiry-based lessons are already available, their adaptation to the knowledge and skill level of the students requires significant effort. Cognitive assistants may help in the generation and the adaptation of generic inquiry-based lessons to the specific instructor needs and learning objectives, by properly instantiating pre-existing patterns.

Teaching inquiry-based lessons requires providing continuous guidance to the students (Kirschner et al., 2006). This is particularly challenging because it cannot be provided globally at the class level (as in a typical classroom), but it needs to be personalized to a team or individual student. While such personalized guidance is feasible for small classrooms, it becomes harder for larger classes. A cognitive assistant can build a model of the student based on their interactions, and automatically update the guidance accordingly. As a result, the instructor needs to offer guidance only for the more subtle issues.

Evaluating students' inquiry seems to be in contradiction with the test-based standards. While inquiry must reward the discovery process followed, the test-based standards reward the right answer and the use of standardized rules. Mitigations of this contradiction exist (e.g., project-based assessment, modifying the grading structure, portfolio-based evaluation, effort, and participation rubrics), but their implementation is difficult and increases the time typically allocated for evaluation. Cognitive assistants used in the process of instruction can collect valuable data for student evaluation (e.g., students' participation in the inquiry process, independence in the inquiry, acquisition of particular skills).

Evaluating the quality and effectiveness of inquiry is challenging because direct student comparisons are not usually possible (Quigley et al., 2011). Cognitive assistants may not only help in monitoring students' performance, but also in offering dynamic feedback to instructors, guiding them to provide better feedback to the students needing it most, and identifying for them the concepts that need reinforced presentation at the end of the classroom.

Scientific Inquiry as Discovery of Evidence, Hypotheses, and Arguments

Our research on developing cognitive assistants for inquiry-based teaching and learning started with defining a computational framework for inquiry, which is abstracted in Figure 1 and explained in the following.

Scientific inquiry is viewed as *ceaseless discovery of evidence, hypotheses, and arguments, through evidence in search of hypotheses, hypotheses in search of evidence, and evidentiary testing of hypotheses.*

Let us assume that we have made a surprising observation of some event in nature (see the bottom left side of Figure 1). The question is: *What hypotheses would explain this observation?* By means of *abductive (imaginative) reasoning*, which shows that something is *possibly* true, we formulate possible hypotheses or explanations. This is what we call "*evidence in search of hypotheses.*"

We would need to analyze each of these hypotheses to determine which of them is the most likely. For this, we need additional evidence which is obtained through the next process called "*hypotheses in search of evidence.*" As shown in the middle part of Figure 1, we put each of the alternative hypotheses to work, generating new lines of inquiry to discover relevant evidence. The question is: *Assuming that this hypothesis is true, what other phenomena should be observable?* Using *deductive reasoning*, which shows that something is *necessarily* true, we decompose our hypothesis into simpler hypotheses and look for evidence that bears upon them. For example, if the hypothesis H would be true, then the subhypotheses H_1 and H_2 would also be true. But if H_1 would be true, then we would need to observe the phenomenon P_1 . We search for it and we may or may not find it.

Some of the newly discovered phenomena or items of evidence may trigger new hypotheses or the refinement of the existing hypotheses. Therefore, as indicated at the bottom left of Figure 1, the processes of evidence in search of hypotheses and hypotheses in search of evidence take place at the same time, and in response to one another.

Now that we have discovered additional phenomena or items of evidence (see the bottom right side of Figure 1) the question is: *What is the probability of the hypothesis based on the available evidence?* Hypothesis testing is necessarily *probabilistic (or inductive)* in nature because our evidence is always *incomplete, usually inconclusive* (it is consistent

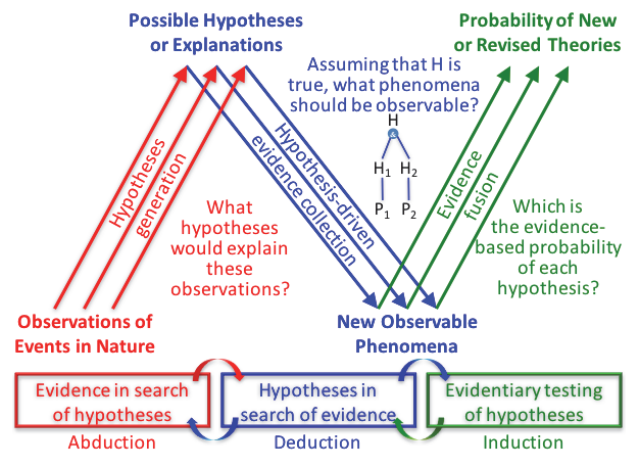


Figure 1. Scientific inquiry as discovery of evidence, hypotheses, and arguments.

with the truth of more than one hypothesis or explanation), frequently *ambiguous* (we cannot always determine exactly what the evidence is telling us), commonly *dissonant* (some of it favors one hypothesis or explanation, but other evidence favors other hypotheses), and with various degrees of *believability* (Schum, 2001; Tecuci et al., 2016). Hypotheses testing depends on the *relevance* and the *believability* of evidence. These factors combine in complex ways to determine the *inferential force* of evidence and the probabilities of the hypotheses, as will be discussed in the next section.

Notice in Figure 1 that the starting point of the inquiry process are observations of events in nature. However, the starting point could also be a question. In such a case the alternative hypotheses to be considered are the possible answers to the question.

Cogent: Cognitive Agent for Cogent Analysis

The framework introduced in the previous section is based on a similar framework that we have previously developed for intelligence analysis. That framework was used as a basis for developing a sequence of increasingly more practical cognitive assistants for the intelligence analysis: Disciple-LTA (Tecuci et al., 2005a; 2007; 2008), TIACRITIS (Boicu et al., 2011; Tecuci et al., 2011), Disciple-CD (Tecuci et al., 2016a), and Cogent (Tecuci et al., 2015).

We are building directly on the latest of these cognitive assistants, Cogent, and on the Disciple learning agents technology (Boicu et al., 2000; Tecuci et al., 2005b; Tecuci et al., 2016b), to develop cognitive assistants for inquiry-based teaching and learning of science, which are described in the following sections.

Inquiry with a Cognitive Assistant in a Fifth Grade Science Classroom

A classical textbook example of using inquiry in a classroom is presented in (NRC, 2000, pp.5-11). In the following, we extend this example with a hypothetical interaction between students and a cognitive assistant, called Inquirer. The goal is to show how such a cognitive assistant can naturally support inquiry-based teaching and learning.

« Several of the students in Mrs. Graham’s fifth grade science class were excited when they returned to their room after recess one fall day. They pulled their teacher over to a window, pointed outside, and said, *we noticed something about the trees on the playground. The left one has lost all its leaves, the middle one has multicolored leaves — mostly yellow — while the right one has lush, green leaves. Why are those trees different? They used to look the same, didn’t they?* Mrs. Graham didn’t know the answer. But she knew that her class was scheduled to study plants later in the year, and this was an opportunity for them

to investigate questions about plant growth that they had originated and thus were especially motivated to answer. Although she was uncertain about where her students’ questions would lead, Mrs. Graham chose to take the risk of letting her students pursue investigations with the Inquirer’s assistance and her guidance. *Let’s make a list of ideas that might explain what’s happening to those trees outside.* They came up with a list of competing explanatory hypotheses, including the following ones (shown also in the current Cogent interface from Figure 2):

- It must be too much water at the root that causes the tree to die.
- The trees have different ages.
- Insects are eating the trees.

She then invited each student to pick one hypothesis which led to several groups, a “water” group, an “age” group, an “illness” group, etc. She asked each group to use the Inquirer assistant in order to plan and conduct a simple investigation to test their preferred hypothesis.

For the next three weeks, science periods were set aside for each group to carry out its investigation. Each group used the Inquirer assistant to conduct its investigation, discovering a variety of sources with information about characteristics of trees, their life cycles, and their environments.

Let us consider the water group that investigated the hypothesis “It must be too much water at the root that causes the tree to die.” Inquirer guided the students to consider how this hypothesis may be decomposed into simpler ones: *Could you think of simpler hypotheses that would support the truthfulness of your hypothesis?* As a result, the students decomposed their hypothesis into the following ones by entering them into Inquirer (see also the top part of Figure 3 showing the hypotheses in the current Cogent interface):

- There is too much water at the root.
- Too much water at the root causes the tree to die.

Then Inquirer instructed the students that they also need to assess the *relevance* or *strength* of these two subhypotheses: *Assuming that these subhypotheses are true, how certain are you that “It is too much water at the root that causes the trees to die?”* The students responded by selecting **certain (C)** from the following list of probabilistic assessments displayed by Inquirer: {**lack of support (LS)**, **likely (L)**, **very likely (VL)**, **almost certain (AC)**, **certain (C)**}.

The students were then instructed by Inquirer to assess the

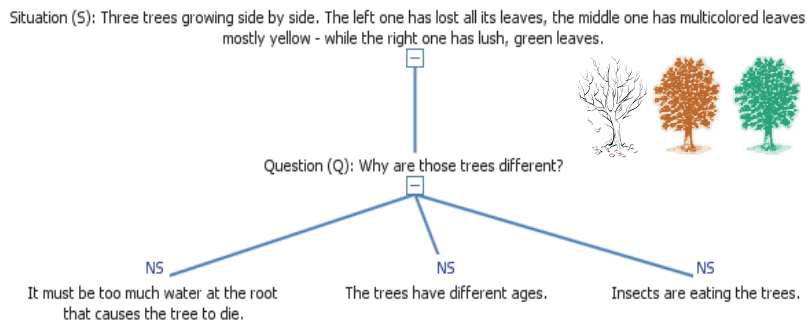


Figure 2: Situation, question, and its possible answers.

two (simpler) subhypotheses based on evidence. To collect evidence for the first one, they decided to look at the ground around the trees every hour that they could. They took turns on making individual observations and since some of them lived near the school, their observations continued after school hours and on weekends. Even though they missed some hourly observations, they had sufficient data indicating that there is too much water at the root of the dying tree, which they introduced into Inquirer as favoring evidence **E1 Water observations** (see the bottom left side of Figure 3).

Next Inquirer guided the students to assess the *believability* and the *relevance* of **E1 Water observations**.

How certain are you that E1 Water observations is true (i.e., that there is indeed too much water at the root of the tree)? The students' answer was **almost certain (AC)** since a few data points were missing and, on rare occasions, the tree was not standing in the water.

Imagine now that you were certain that E1 Water observations is true. In such a case, how certain would you be that there is too much water at the root of the tree? **certain (C)**.

Based on the students' assessments, Inquirer determined that the *inferential force* of **E1 Water observations** on the hypothesis "There is too much water at the root" is **almost certain (AC)**, as shown in Figure 3. Inquirer explained to the students that *inferential force* answers the question: *How strong is E1 Water observations in favoring the hypothesis "There is too much water at the root"?* An item of evidence will convince us that a hypothesis is true if and only if the item is both highly believable and highly relevant. Therefore, the inferential force was computed as the minimum of the believability of **E1 Water observations** (**almost certain - AC**) and its relevance of (**certain - C**), minimum which is **almost certain**. Because **E1 Water observations** is currently the only item of evidence relevant to the hypothesis "There is too much water at the root," the probability of this hypothesis is also **almost certain (AC)**.

Then, one of the students recalled that several months ago the leaves on one of his mother's geraniums had begun to turn yellow. She told him that the geranium was getting too much water. This item of information was represented in Inquirer as item of evidence **E2 Geranium case** favoring the hypothesis "Too much water at the root causes the tree to die." The students agreed to assess its believability as **almost certain** (because, although the mother has experience with plants, she is not a professional), and its relevance as **very likely** (because geraniums is a different type of plant), leading Inquirer to compute its inferential force as **very likely**.

Additionally, Mrs. Graham gave the group a pamphlet from a local nursery entitled "Growing Healthy Plants." The water group read the pamphlet and found that when plant roots are surrounded by water, they cannot take in air from the space around the roots and they essentially "drown." This item of information was represented in Inquirer as an additional item of evidence **E3 Growing Healthy Plants** favoring the hypothesis "Too much water at the root causes

the tree to die." The students agreed to assess its *believability* as **certain** (because this information is from a highly reputable source), and its relevance as **certain** (since it, in fact, asserted the hypothesis), leading Inquirer to compute its inferential force as **certain**. Additionally, Inquirer computed the inferential force of all favoring evidence (i.e., both **E2 Geranium case** and **E3 Growing Healthy Plants**) as **certain**, by taking the maximum of their inferential forces. This is also the probability of the hypothesis "Too much water at the root causes the tree to die" because no disfavoring evidence was found. However, if any disfavoring evidence would have been found, then Inquirer would have determined whether, on balance, the totality of evidence favors or disfavors the hypothesis, and to what degree.

Having assessed the probability of "There is too much water at the root" as **almost certain**, and that of "Too much water at the root causes the tree to die" as **certain**, Inquirer inferred that the probability of their top-level hypothesis "It must be too much water at the root that causes the tree to die" as **almost certain**. This is the minimum between these probabilities and the joint relevance of the two subhypotheses, which is **certain** (see the top part of Figure 3).

Finally, Inquirer automatically generated a report describing the analysis logic, citing sources of data used, and the manner in which the analysis was performed. The report was further edited by the water group before being presented to the class, together with the reports of the other teams.

As different groups presented and compared their analyses, the class learned that some evidence — such as that from the group investigating whether the trees have different ages — did not explain the observations. The results of other investigations, such as the idea that the trees could have a disease, partly supported the observations. But the explanation that seemed most reasonable to the students, that fit all the observations and conformed with what they had learned from other sources, was "too much water." After their three weeks of work, the class was satisfied that together they have found a reasonable answer to their question. » (adapted from NRC, 2000, pp.5-11).

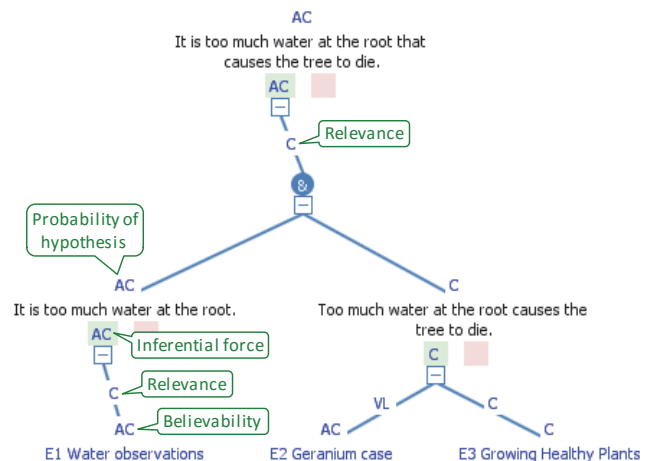


Figure 3: Hypothesis analysis.

Teaching Critical Thinking Skills in Science with *sInvestigator*

We have started the design and development of a cognitive assistant, called *sInvestigator*, that will help undergraduate students develop critical thinking skills in addressing scientific problems. *sInvestigator* implements the computational model of inquiry presented and illustrated in the previous sections, in the form of general inquiry methods and general knowledge about the properties, uses, discovery, and marshaling of evidence. This enables *sInvestigator* both to teach the students how to approach complex scientific problems, and to assist them in actually solving new problems.

Such an engagement of learners in an inquiry process to generate hypotheses, gather data, and evaluate evidence is aligned with constructivist learning theories, particularly those related to discovery learning and problem-solving (Bruner, 1961). Inquiry-based science teaching is commonly advocated (NRC, 1996, 2000; NGSS Lead States, 2013), but the use of inquiry methods alone may not be inherently engaging (Hampden-Thompson and Bennett, 2013). Students using *sInvestigator* will use and evaluate evidence and hypotheses within the relevant course content. This will provide a meaningful context within which students may develop a deeper understanding of critical scientific practices.

The next sections illustrate the planned use of *sInvestigator* in the GMU undergraduate course HNRS 240 History of Science. The actual example concerns the adoption of the theory of plate tectonics in the 1950's and 60's. The goal is to teach the students (sophomores from all disciplines) both critical scientific practices and content knowledge.

Hypotheses Generation through Evidence in Search of Hypotheses

At the beginning of the class the students are guided in the abductive process of hypotheses discovery represented in the left side of Figure 1. The instructor presents the students a world map in the evidence module of *sInvestigator* (see Figure 4), pointing to the shapes of the continents, and asking them to propose some hypotheses or explanations of these shapes: “Q: How are the shapes of the continents?” (see bottom left side of Figure 5). If the students do not propose any hypothesis, the instructor will give them more detailed guidance

(e.g., pointing specifically to the boundaries of Africa and South America), leading them to propose as alternative answers, the two hypotheses from the bottom of Figure 5.

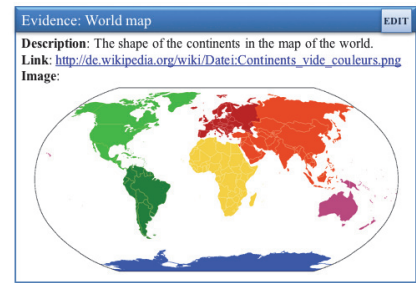


Figure 4: A map of the world.

The instructor continues to guide the students in asking questions that explore possible higher-level explanations of the World map evidence (see the left hand side of Figure 5), until they develop the abductive reasoning chain from the middle of Figure 5. The top hypothesis in the chain will be obtained by asking the students to investigate which is the current theory related to the movement of the continents.

Notice that each question asked has, in fact, multiple possible answers (see the hypotheses from the middle and right side of Figure 5). In order to conclude that the top hypothesis “H: The continents are the land part of plates that make up the surface of the earth, and the plates are continuously moving” is true, one would need to show that each hypothesis on the chain from the evidence “E1: World map” to this top hypothesis is more likely than its alternative hypothesis. For this, however, one needs evidence to test all these competing hypotheses. The next section explains how such evidence can be obtained using *sInvestigator*'s guidance.

Evidence Discovery through Hypotheses in Search of Evidence

As indicated in the middle of Figure 1, *sInvestigator* guides the students to put each of the hypotheses from Figure 5 to work to generate new lines of inquiry and obtain new evidence. The strategy is to decompose each such hypothesis, starting with the simpler hypotheses from the bottom of the

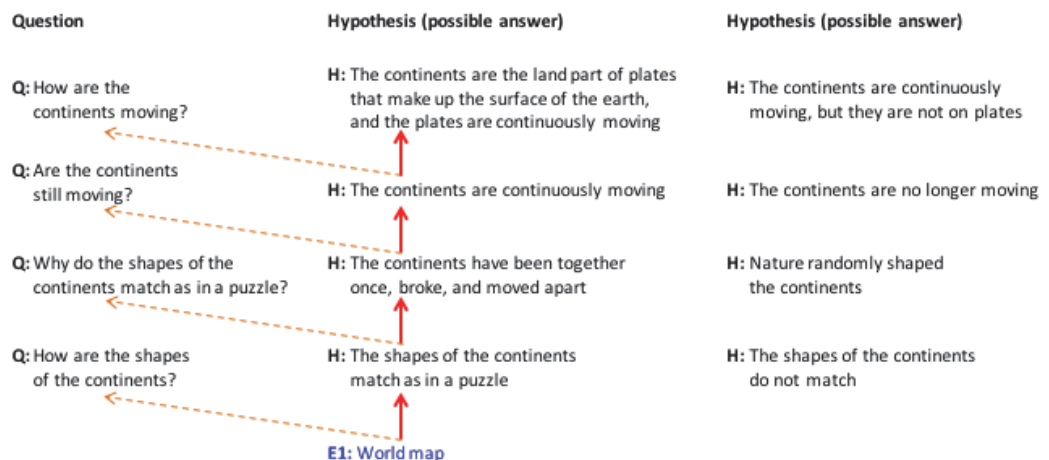


Figure 5. Discovery of hypotheses.

figure.

For example, Figure 6 shows the decomposition of the hypothesis “**H**: The continents have been together once, broke, and moved apart” into a conjunction of two subhypotheses. Each of them is further decomposed into disjunctions of even simpler hypotheses that show much more clearly what evidence to search for. The students will need to look for evidence that either favors or disfavors the leaf hypotheses in this tree, and we expect them to find evidence such as that from the bottom left side of Figure 6:

E2: Similar shapes of continental shelf, another evidence (in addition to **E1**) supporting the hypothesis “**H**: The shapes of the continents match as in a puzzle.”

E3: Identical fossils on both sides of the Atlantic, supporting the hypothesis “**H**: There are continuous fossil records across continents.”

E4: Tropical fossils in Antarctica, supporting the hypothesis “**H**: There are fossils records suggesting different continent positions in the past.”

The other hypotheses from Figure 5 are decomposed in a similar manner to guide the search for evidence.

Evidentiary Testing of Hypotheses

Once evidence for a hypothesis is found, it can be used to assess the probability of that hypothesis. This corresponds to the evidentiary testing of hypotheses phase from the right hand side of Figure 1.

For example, Figure 7 shows the probabilistic assessment of the hypothesis “**H**: The continents are continuously moving.” To obtain it, the students have first to assess the believability of each item of evidence and its relevance to the corresponding elementary hypothesis, leading to the probability of each elementary hypothesis. Then, the probabilities

of the upper level hypotheses are obtained from the probabilities of the lower level hypotheses, by following the logic of the argumentation structure.

As discussed previously and illustrated in Figure 5, students are guided to consider more and more complex hypotheses that explain an observation, such as the shapes of the continents in Figure 4.

They first collect evidence for the simplest competing hypotheses from the bottom of Figure 5.

Having assessed that the most likely of these two hypotheses is “**H**: The shapes of the continents match as in a puzzle,” they continue with searching evidence for the next upper level competing hypotheses in Figure 5, and with testing them, concluding that the most likely of these hypotheses is “**H**: The continents have been together once, broke, and moved apart.”

This process continues until the probabilities of the top level hypotheses in Figure 5 are assessed. At this point the students have succeeded in providing evidential support to the plate tectonics theory.

The students will work in teams, each team developing its own logic. Then they will present and debate the developed argumentations and evidence, and will work together to develop a better argumentation that integrates their best ideas and evidence.

Types of Exercises with *sInvestigator*

While the previous section has illustrated the envisioned use of *sInvestigator* to completely answer a scientific question, the system could also be used to introduce various concepts and methods through shorter exercises.

In one exercise, *sInvestigator* will teach the students how a complex hypothesis is decomposed into simpler hypotheses, and how the assessments of these simpler hypotheses

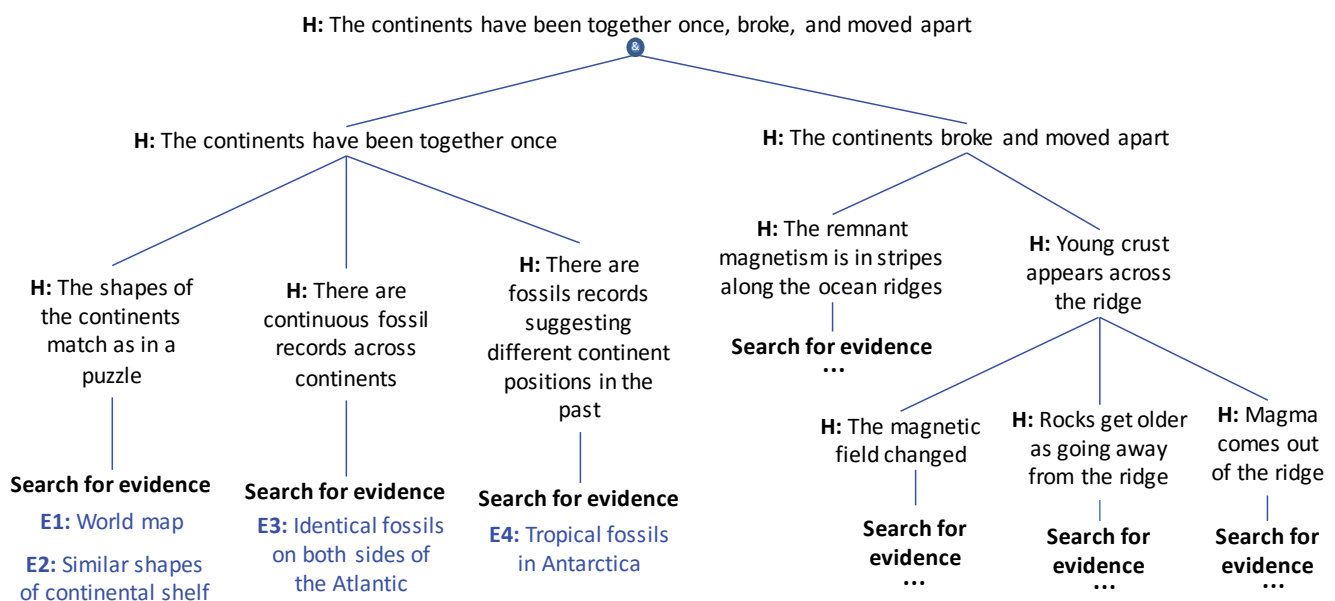


Figure 6. Hypotheses in search of evidence.

are combined into the assessment of the more complex hypothesis (as illustrated in Figure 7).

In another exercise, *sInvestigator* will teach the students how to assess the relevance and the believability of evidence. This exercise will provide both a decomposition tree and a set of items of information, some relevant to the considered hypotheses and some irrelevant. The students will be asked to determine which item of information is relevant to which hypothesis, and whether it is favoring or disfavoring evidence. They will also be asked to assess and justify the relevance and the believability of each item of evidence. After completing their analyses, the teacher and *sInvestigator* will provide additional information, asking the students to update their analyses in the light of the new evidence. Finally, the students will present, compare, and debate their analyses in class.

In another exercise, *sInvestigator* will provide an analysis tree but no items of evidence, asking the students to look for relevant evidence (e.g., by searching the Internet, or by performing various experiments), and complete the analysis.

In a more complex exercise, *sInvestigator* will present a scenario with a surprising observation, like the shape of the continents discussed in the previous section. The students

will be asked to formulate competing hypotheses that may explain the surprising observation, use the formulated hypotheses to discover evidence, and use the discovered evidence to assess each hypothesis. Then they will compare their analyses of the competing hypotheses, and will select the most likely hypothesis. *sInvestigator* will assist the students in this process by guiding them in decomposing hypotheses, in searching for evidence, in assessing the hypotheses, in combining the assessments of the simpler hypotheses, and in comparing the analyses of the competing hypotheses.

sInvestigator will provide many opportunities for collaborative work. For example, a complex hypothesis may be decomposed into simpler hypotheses, each assessed by a different student. Then the results obtained by different students will be combined to produce the assessment of the complex hypothesis. Or different students will analyze the same hypothesis. Then they will compare and debate their analyses and evidence, and work together toward producing a consensus analysis.

sInvestigator is envisioned as both a teaching tool for teachers, and as a learning assistant for students. For example, the teacher will demonstrate some of these exercises in

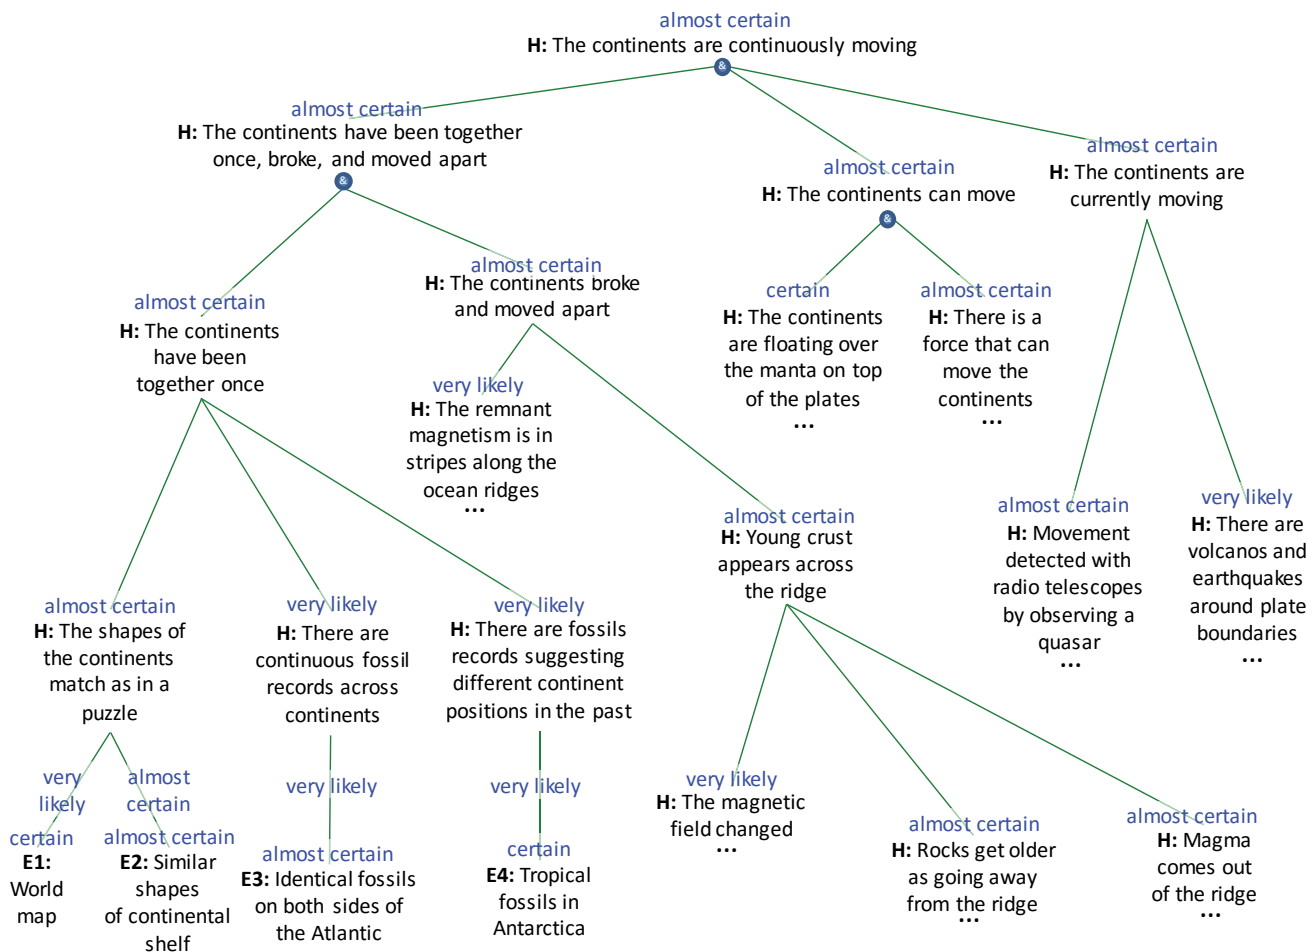


Figure 7. Evidence-based hypothesis assessment.

class. Other exercises will be performed by the students, under the guidance of the teacher, and with the assistance of *sInvestigator*.

Exercises like the ones described above will make more clear what is the level of understanding of each student for each of the major aspects of the scientific inquiry. Thus *sInvestigator* will also facilitate a more objective assessment.

sInvestigator may also be used to train teachers for inquiry-based teaching and learning. This can be done in a classroom setting. But it can also be done through individual learning because *sInvestigator* can be installed on teachers' own computers. Moreover, *sInvestigator* will come with a stock of exercises like those discussed above.

Conclusions and Future Work

This paper has presented current research on designing and developing cognitive assistants for supporting inquiry-based teaching and learning of science.

Our future work will focus on the further development, experimental use, and evaluation of *sInvestigator* in two undergraduate honors courses at George Mason University, HNRS 353 Technology in the Contemporary World (Modern and Scientific Revolution) and HNRS 240 Reading the Past (History of Science), addressing various topics, such as the atomic structure, the Copernican revolution, and the theory of evolution, in addition to the plate tectonics theory. After that, we plan to introduce *sInvestigator* to the general student body in the course PROV 301 Great Ideas in Science, which is open to any George Mason student.

In the longer term we plan to investigate the development and use of such cognitive assistants in other disciplines, and not only for undergraduate students, but also for K-12.

References

Boicu, M., Tecuci, G., Marcu, D., Bowman, M., Shyr, P., Ciucu, F., and Levcovici, C. 2000. Disciple-COA: From Agent Programming to Agent Teaching, *Proc. 27th Int. Conference on Machine Learning (ICML)*, Stanford, California, Morgan Kaufman.

Boicu, M., Marcu, D., Tecuci, G., Schum, D.A. 2011. Cognitive Assistants for Evidence-Based Reasoning Tasks, *AAAI Fall Symposium on Advances in Cognitive Systems*, Arlington, VA, 4-6 November.

Bransford, J. D., and Donovan, M. S. eds. 2004. *How Students Learn: History in the Classroom*. National Academies Press.

Brown, P. L., Abell, S. K., Demir, A., and Schmidt, F. J. 2006. College science teachers views of classroom inquiry. *Science Education*, 90(5), 784-802.

Bruner, J. S. 1961. The act of discovery. *Harvard Educational Review*, 31, 21-32.

Hampden-Thompson, G., and Bennett, J. 2013. Science teaching and learning activities and students' engagement in science. *International Journal of Science Education*, 35(8), 1325-1343.

Kirschner, P.A., Sweller, J., Clark, R.E. 2006. Why Minimal Guid-

ance During Instruction Does Not Work: An Analysis of the Failure of Constructivist Discovery, Problem-based, Experiential, and Inquiry-Based Teaching. *Educational Psychologist*, 41(2), 75-86, Lawrence Erlbaum Associates, Inc.

Quigley, C., Marshall, J. C., Deaton, C. C. M., Cook, M., Padilla, M. 2011. Challenges to Inquiry Teaching and Suggestions for How to Meet Them. *Educational Journal*, 20(1), 55-61.

Minner, D. D., Levy, A. J., and Century, J. 2010. Inquiry-based science instruction—what is it and does it matter? Results from a research synthesis years 1984 to 2002. *Journal of Research in Science Teaching*, 47(4), 474-496.

Oliver-Hoyo, M., Anderson, M., and Allen, D. 2004. Inquiry-guided instruction: practical issues of implementation. *Journal of College Science Teaching*, 33(6), 20.

NGSS Lead States. 2013. *Next Generation Science Standards: For States, By States*. The National Academies Press.

NRC 1996. National Research Council: *National Science Education Standards*. Washington, DC: National Academy Press.

NRC 2000. National Research Council: *Inquiry and the National Science Education Standards* Washington, DC: National Academy Press.

POGIL 2016, Process-Oriented Guided-Inquiry Learning, <https://pogil.org/>, <https://pogil.org/resources/curriculum-materials>

Schum, D. A. 2001. *The Evidential Foundations of Probabilistic Reasoning*. Northwestern University Press.

Tecuci, G., Boicu, M., Ayers, C., and Cammons D. 2005a. Personal Cognitive Assistants for Military Intelligence Analysis: Mixed-Initiative Learning, Tutoring, and Problem Solving. In *Proc. 1st Int. Conf. on Intelligence Analysis*, McLean, VA, 2-6 May.

Tecuci, G., Boicu, M., Boicu, C., Marcu, D., Stanescu, B., Barbulescu, M. 2005b. The Disciple-RKF Learning and Reasoning Agent, *Computational Intelligence*, 21(4), 462-479.

Tecuci, G., Boicu, M., Marcu, D., Boicu, C., Barbulescu, M., Ayers, C., and Cammons D. 2007. Cognitive Assistants for Analysts, *Journal of Intelligence Community Research and Development (JICRD)*. Also in Auger J. and Wimbish W. eds. *Proteus Futures Digest*, 303-329, Joint publication of National Intelligence University, Office of the Director of National Intelligence, and US Army War College Center for Strategic Leadership.

Tecuci, G., Boicu, M., Marcu, D., Boicu, C., Barbulescu, M. 2008. Disciple-LTA: Learning, Tutoring and Analytic Assistance, *Journal of Intelligence Community Research and Development*, July.

Tecuci, G., Marcu, D., Boicu, M., Schum, D. A., and Russell, K. 2011. Computational Theory and Cognitive Assistant for Intelligence Analysis, in *Proceedings of the Sixth International Conference on Semantic Technologies for Intelligence, Defense, and Security – STIDS*, 68-75, Fairfax, VA, 16-18 November.

Tecuci, G., Marcu, D., Boicu, M., Schum, D. A. 2015 COGENT: Cognitive Agent for Cogent Analysis, in *Proceedings of the 2015 AAAI Fall Symposium "Cognitive Assistance in Government and Public Sector Applications"*, 58-65, Arlington, VA, 12-14 November, Technical Report FS-15-02, AAAI Press, Palo Alto, CA.

Tecuci, G., Schum, D. A., Marcu, D., and Boicu, M. 2016a. *Intelligence Analysis as Discovery of Evidence, Hypotheses, and Arguments: Connecting the Dots*, Cambridge University Press.

Tecuci G., Marcu D., Boicu M., Schum D.A. 2016b. *Knowledge Engineering: Building Cognitive Assistants for Evidence-based Reasoning*, Cambridge University Press.