18th International Conference on Control Systems and Computer Science,
University Politehnica of Bucharest, May 24-27, 2011

Toward a Computational Theory of Evidence-based Reasoning

Gheorghe Tecuci*, David Schum, Mihai Boicu***, Dorin Marcu, Katherine Russell**

Learning Agents Center, The Volgenau School of Engineering, George Mason University, VA 22030, USA
{tecuci, dschum, mboicu, dmarcu, krusselc}@gmu.edu
*Also with * Department of Computer Science, ** Department Systems Engineering and Operations Research and the School of Law, *** Department of Applied Information Technology*

Abstract: Evidence-based reasoning is at the core of many problem solving and decision making tasks in a wide variety of domains. Through abductive reasoning we generate hypotheses from our observations; through deductive reasoning we use our hypotheses to generate new lines of inquiry and discover new evidence; and through inductive reasoning we test our hypotheses with this discovered evidence. These processes, which integrate imaginative and critical reasoning, are often stunningly complex because our evidence is incomplete, inconclusive, ambiguous, dissonant, and has various degrees of believability. We present research performed in the Learning Agents Center of George Mason University on developing a Computational Theory of Evidence-based Reasoning viewed as mixed-initiative integration of evidence in search of hypotheses, hypotheses in search of evidence, and evidential tests of hypotheses, all taking place simultaneously, in a world that is changing all the time. This theory is embedded in the Disciple cognitive agents which are capable of capturing tacit knowledge from subject matter experts, and can act as assistants to experts, as expert consultants to non-experts, or as intelligent tutors to students. We illustrate the applications of these agents in various domains, including intelligence analysis and inquiry-based learning in natural sciences.

1. INTRODUCTION

“In its simplest sense, evidence may be defined as any factual datum which in some manner assists in drawing conclusions, either favorable or unfavorable, to some hypothesis whose proof or refutation is being attempted.” (Murphy, 2003, p.1).

Evidence-based reasoning is at the core of many problem solving and decision making tasks in a wide variety of domains, including law, intelligence analysis, forensics, medicine, physics, chemistry, history, archaeology, and many others. This is not surprising because, as Jeremy Betham stated over two centuries ago, “The field of evidence is no other than the field of knowledge” (Betham, 1810). What is surprising is that evidence-based reasoning processes in these domains have been studied in isolation from each other.

It is only recently that steps have been taken to study evidence as a multidisciplinary subject (Twinning, 2003), and to lay the foundation for a *Science of Evidence* (Schum, 2009). This paper continues this trend by suggesting the development of a general *Computational Theory of Evidence-based Reasoning*, and presenting some recent results obtained in the Learning Agents Center of George Mason University. The immediate question that one could ask is: Why a computational theory? The answer is because evidentiary reasoning is frequently of such an astonishing complexity that it can be best approached through the mixed-initiative integration of human imagination and computer knowledge-based reasoning (Tecuci et al., 2007). A main reason for this complexity is that our evidence is always *incomplete* (we can look for more, if we have time), usually *inconclusive* (it is consistent with the truth of more than one hypothesis or possible explanation), frequently *ambiguous* (we cannot always determine exactly what the evidence is telling us), commonly *dissonant* (some of it favors one hypothesis or possible explanation but other evidence favors other hypotheses), and has various degrees of *believability* (Schum, 2001a; Tecuci et al., 2010b). Arguments, requiring both *imaginative and critical reasoning* and involving all known types of inference (*deduction*, *induction*, and *abduction*), are necessary in order to establish and defend the three major credentials of evidence: its *relevance*, its *believability*, and its *inferential force or weight* with respect to the considered hypotheses.

The next section presents a general view on evidence-based reasoning that summarizes many thoughts expressed over the centuries, showing how it applies to several domains. After that, Section 3 presents in more detail an example of evidence-based hypothesis generation and analysis in the intelligence domain. Section 4 provides some details on the implementation of the computational theory of evidence-based reasoning in the Disciple cognitive agents. Section 5 presents how evidence-based reasoning applies to inquiry-based learning in natural sciences. Section 6 discusses five complexity issues in evidence-based reasoning and how the proposed approach attempts to cope with them. Finally, the conclusion of this paper argues for studying evidence-based reasoning as a general critical thinking skill, at all levels of education, from the elementary school to the university, proposing Disciple-based agents as educational tools for learning complex reasoning skills through a hands-on approach.

2. EVIDENCE, HYPOTHESES, AND ARGUMENTS

Evidence-based reasoning involves the discovery of evidence, hypotheses, and arguments linking them (Schum, 1987; 2001a). Ever since Aristotle (384BC-322BC), some of the greatest minds have been interested in the process of discovery and, in particular, in understanding the distinction between discovering a hypothesis and testing it. Galileo (1564-1642) thought that we "reason backward" inductively to imagine causes (hypotheses) from observed events, and we reason deductively to test the hypotheses. A similar view was held by Isaac Newton (1642-1727), John Locke (1632-1704), and William Whewell (1794-1866). Charles S. Peirce (1839-1914) was the first to suggest that new ideas or hypotheses are generated through a different form of reasoning, which he called *abduction* and associated with imaginative reasoning (Peirce, 1898; 1901). His views are very similar to those of Sherlock Holmes, the famous fictional character of Conan Doyle.

Fig. 1 summarizes our view on the discovery of evidence, hypotheses, and arguments, which builds on further study of these processes by many people, especially Wigmore (1913; 1937) and Schum (1987; 2001a).

Let us suppose, as in Science, that we already have a collection of prior evidence in some investigation and an existing collection of hypotheses H_1, H_2, … , H_n, which explains this prior evidence. But now we make an observation E* which is not explained by any of these n hypotheses. The question is: *What other hypothesis would explain this observation?* Through *abductive (imaginative) reasoning*, which shows that something is *possibly* true, we generate the hypothesis H_{n+1}. This may appear to us as a "flash of insight," sometimes much later, when we are occupied with other things, and it may not be immediately clear to us why H_{n+1} is possible. However, further thoughts may produce a chain of reasoning like: E* → it is possible that F is true → it is possible G is true → it is possible that H_{n+1} is true. Notice that this is a process of *evidence in search of hypotheses*, where we look for hypotheses that explain our observations (see the left hand side of Fig. 1).

But no matter how novel or imaginative our new hypothesis H_{n+1} is, it would not be very appealing if it would only explain E*. What we would like for H_{n+1} is to explain our prior observations better than the previously generated hypotheses, and to also suggest new potentially observable evidence that our previous hypotheses did not suggest. So we put this hypothesis at work by asking the question: *Assuming that H_{n+1} is true, what other things should be observable?* Then, through *deductive reasoning* which shows that something is *necessarily* true, we successively determine other hypotheses that would need to be true if H_{n+1} were true, and the observations that they would entail. This opens new lines of inquiry and guides us to identify new items of evidence. Notice that this is a process of *hypotheses in search of evidence*, where we look for evidence that is entailed by our hypotheses (see the middle of Fig. 1).

Now, some of the newly discovered items of evidence may trigger new hypotheses, or the refinement of the current hypothesis. Therefore, as indicated at the bottom left of Fig. 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. They result in hypotheses that have to be tested, through *inductive reasoning*, which shows that something is *probably* true (see the right hand side of Fig. 1). Such testing depends on the relevance and believability of our evidence. These factors combine in further complex ways in argumentation structures that allow us to assess the inferential force or weight of the evidence we are considering.

We think that the same type of evidence-based reasoning occurs in many domains. Scientists from various domains, such as physics, chemistry, or biology, may recognize this as a formulation of the scientific method (Noon, 2009).

In law, an attorney makes observations in a criminal case and seeks to generate hypotheses in the form of charges that seem possible in explaining these observations. Then, assuming that a charge is justified, attempts are made to deduce further evidence bearing on it. Finally, the obtained evidence is used to prove the charge.

In medicine, a doctor makes observations with respect to a patient's complaints and attempts to generate possible diagnoses (hypotheses) that would explain them. She then performs various medical tests that provide further evidence which is used in forming a final diagnosis for the patient.

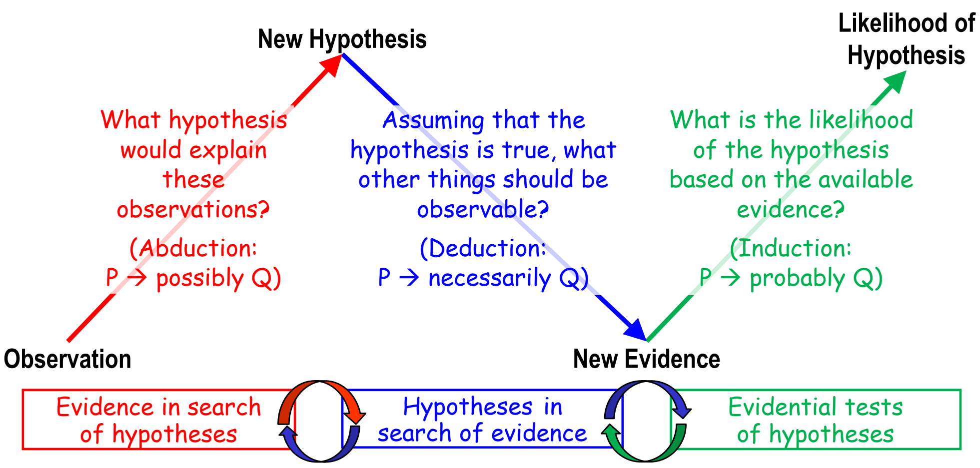

Fig. 1. Discovery of evidence, hypotheses, and arguments

In forensics, observations made at the site of an explosion in a power plant lead to the formulation of several possible causes. Analysis of each possible cause leads to the discovery of new evidence that eliminates or refines some of the causes, and may even suggest new ones. This cycle continues until enough evidence is found to determine the most likely cause.

The next section provides a more detailed illustration of this process in intelligence analysis.

3. INTELLIGENCE ANALYSIS

The objective of intelligence analysis is to answer complex questions arising in the problem solving, planning, or decision-making process, such as, "Does Al Qaeda have nuclear weapons?" or "Will the United States be the world leader in alternative fuels within the next decade?"

Let us consider an intelligence analyst whose mission is to discover threats to the NATO forces in Afghanistan, such as ambushes, kidnappings, rocket launches, false check-points, suicide bombers, or improvised explosive devices. Browsing wide area motion imagery (WAMI) of a region around Kabul, she notices evidence of road work at 1:17AM, at a highway junction located about 11 km northeast of Kabul. Road work at this location at this time is very unusual and the analyst wonders whether it may suggest any threat. Through a flash of insight, she abductively leaps to the hypothesis H_k that there is an ambush threat at that location. Attempting to justify this hypothesis, she generates the following abductive inference steps, shown in the left hand side of Fig. 2: "There is evidence of road work → It is possible that there is indeed road work → It is possible that the road work is for blocking the road → It is possible that this is ambush preparation → It is possible that there is an ambush threat." While this chain of reasoning shows the potential of an ambush threat, there are also alternative hypotheses to be considered. For example, just because we have evidence E^*_i about the event E_i does not mean that E_i actually occurred. At issue here is the *believability* of E^*_i, as will be discussed later. Thus, an alternative hypothesis is always "$\neg E_i$". Similarly, for each of the other hypotheses in the chain from E^*_i to H_k (e.g., H_a: Road blocking) there will generally be alternative hypotheses (e.g., H_{a1}: Road repair). This is the *evidence in search of hypotheses* process where newly discovered evidence searches for hypotheses that would explain it.

Then we put each of these hypotheses at work, starting from bottom-up, to generate new lines of inquiry and evidence. By means of deductive reasoning we decompose our hypothesis into simpler hypotheses and look for evidence that bears upon them. The middle of Fig. 2 illustrates this process for the top-level hypothesis H_k: "If there is indeed an ambush threat at this location then the location must be appropriate for ambush, there should be evidence of ambush preparation in the motion imagery, and (depending on the time of observation) there should be evidence of ambush execution. Furthermore, to be an appropriate ambush location, it should be on a route used by the NATO forces, and there should be cover at that location, etc." As one can notice, each of these sub-hypotheses allows the analyst to deduce potential items of evidence (shown as shaded circles) that bear upon them. So here we have *hypotheses in search of evidence* that may favor or disfavor them.

Collected evidence is then used to test, through *inductive reasoning*, each hypothesis, as illustrated in the right-hand side of Fig. 2 with H_k. This is a *Wigmorean probabilistic inference network* (Wigmore, 1913; 1937; Schum, 2001a) that shows how the available items of evidence (shown as black circles) are linked to the intermediary hypotheses and to H_k through an *argument* that establishes and fuses the *relevance*, *believability* and *inferential force* of a wide variety of evidence of different types (e.g., HUMINT, IMINT, COMINT, SIGINT, MASINT, and Open Source). The result is the likelihood of H_k (It is very likely that there is an ambush threat at location L_1 after 1:17am).

As shown above, intelligence analysis may be understood as

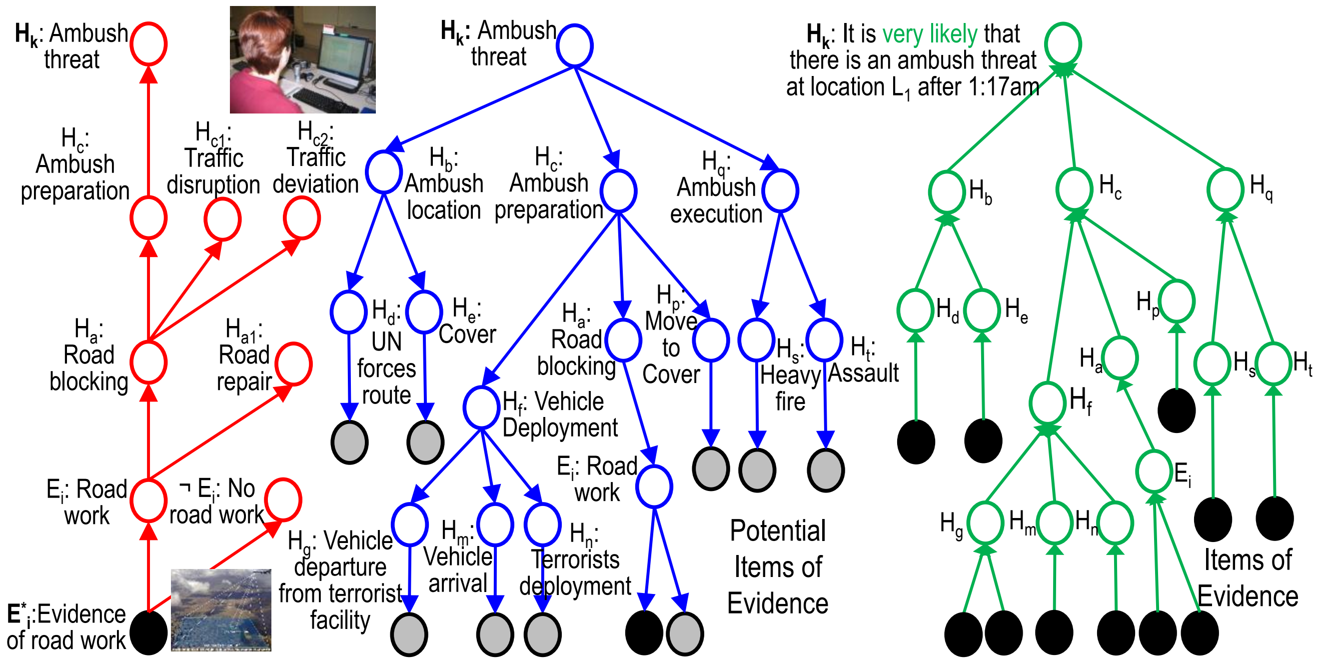

Fig. 2. Hypothesis generation and analysis in Intelligence Analysis

ceaseless discovery of evidence, hypotheses, and arguments, where continuous processes of evidence in search of hypotheses, hypotheses in search of evidence, and evidentiary tests of hypotheses take place at the same time and in response to one another, in a world that is changing all the time, while we are trying to understand it.

4. DISCIPLE COGNITIVE ASSISTANTS FOR INTELLIGENCE ANALYSIS

The proposed computational theory of evidence-based reasoning is partially implemented in Disciple cognitive assistants for intelligence analysts (Tecuci et al., 2008; 2010a, b). These agents synergistically integrate three complex capabilities: (1) they can rapidly learn the analytic expertise which takes years to establish, is lost when analysts separate from service, and is costly to replace; (2) they can tutor new intelligence analysts on how to systematically analyse complex hypotheses; and (3) they can assist the analysts in evaluating the likelihood of hypotheses.

An expert analyst can teach a Disciple agent in a way that is similar to how she would teach a student, through problems solving examples and explanations. She would show it, for instance, reasoning trees like the ones in Fig. 2, and the agent would learn general reasoning rules from the corresponding reasoning steps. Currently, a Disciple agent can learn deductive and inductive rules, and research is being performed to also learn abductive rules (Tecuci et al., 2008). An important issue, however, is that a Disciple agent does not start the learning process with a blank mental tablet, because it already has a stock of established knowledge about evidence, such as its properties, uses, and discovery (Schum, 2001a).

In this paper we will only briefly present the process of hypothesis analysis with these agents, to provide additional details on the actual implementation of the proposed computational theory of evidence-based reasoning.

Fig. 3 shows how the generated hypotheses are assessed by employing a divide and conquer approach (called *problem reduction and solution synthesis*) which combines the deductive and inductive reasoning trees from Fig. 2. This approach is grounded in the problem reduction representations developed in the field of artificial intelligence (Nilsson, 1971; Tecuci, 1988, 1998), and in the argument construction methods provided by the noted jurist John H. Wigmore (1937), the philosopher of science Stephen Toulmin (1963), and the evidence professor David Schum (1987, 2001a). In this approach: (1) the problem of assessing a complex hypothesis is successively reduced to the assessment of simpler and simpler hypotheses; (2) the simplest (elementary) hypotheses are assessed based on the available evidence; and finally, (3) the solutions of these assessments are successively combined, from bottom-up, to obtain the solution of the top level hypothesis assessment.

While the tree in the middle of Fig. 2 shows how the assessment of the hypothesis H_k is reduced to the assessment of simpler and simpler hypotheses, the tree in Fig. 3 shows how the elementary hypothesis H_e (there is cover at location L_1) is assessed based on the available evidence. As indicated, one has to consider both *favoring evidence* and *disfavoring evidence*. In this example there is one item of favoring evidence, E_l (Wide area motion imagery evidence of brush and trees, as well as ruined buildings, which could provide cover). Therefore one has to assess to what extent this item favors the hypothesis H_e. This requires assessing the *relevance* and the *believability* of E_l, and its *inferential force* on H_e.

Relevance answers the question: “So what? How does this item of evidence bear on what the agent is trying to prove or disprove?” *Believability* answers the question: “Can the agent believe what this item of evidence is telling it?” *Inferential force* answers the question: “How strong is this item of relevant evidence in favoring or disfavoring various alternative hypotheses or possible conclusions being entertained?”

The assessment of the believability of E_l is further decomposed into the assessment of its believability credentials. E_l is an image of tangible things (i.e.

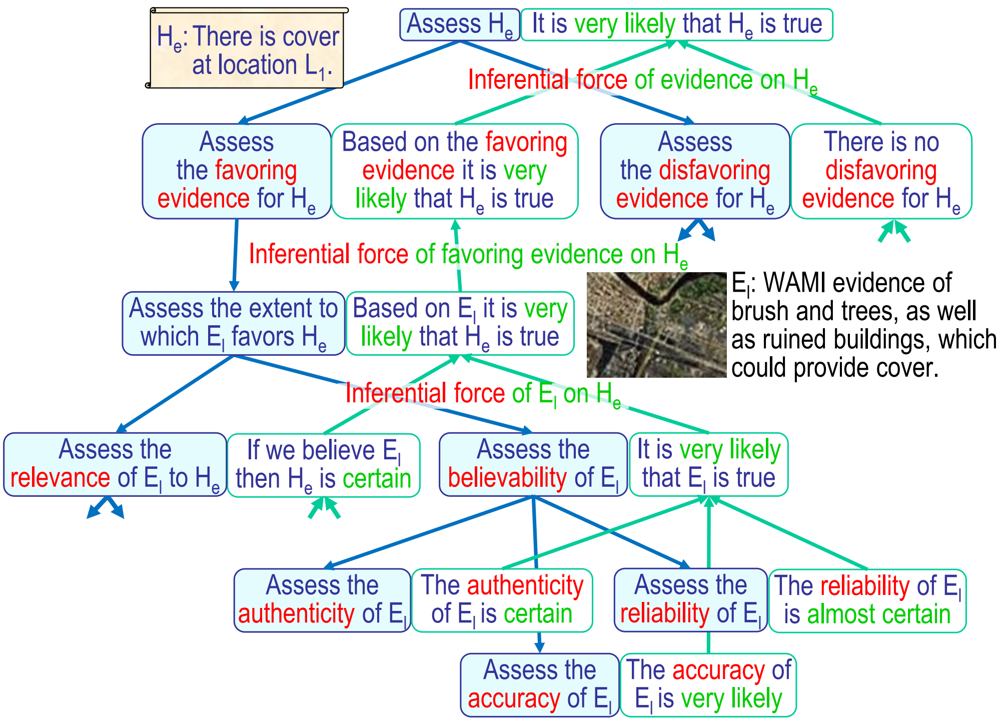

Fig. 3. Hypothesis assessment through problem reduction and solution synthesis

demonstrative tangible evidence). Therefore it has three believability attributes, as shown at the bottom of Fig. 3. The first believability attribute is *authenticity* which addresses the question: "Is this object what it is represented as being or is claimed to be?" The second believability attribute is *accuracy* of the representation provided by the demonstrative tangible item. The accuracy question concerns the extent to which the device that produced the representation of the real tangible item had a degree of sensitivity (resolving power or accuracy) that allows us to tell what events were observed. The third major attribute, *reliability*, is especially relevant to various forms of sensors that provide us with many forms of demonstrative tangible evidence. A system, sensor, or test of any kind is reliable to the extent that the results it provides are repeatable or consistent; one says that a sensing device is reliable if it would provide the same image or report on successive occasions on which this device is used.

In the illustration from Fig. 3 we assumed the following fuzzy assessments of these three attributes: "certain," "very likely," and "almost certain." They are combined, through a "min" function, to produce the following assessment of the believability of E_l: "It is very likely that E_l is true." Then, by combining this assessment with the conditional relevance of E_l ("If we believe E_l then H_e is almost certain"), again through a "min" function, one obtains the inferential force of E_l on H_e ("Based on E_l it is very likely that H_e is true"). Similarly the agent would assess the inferential force of other items of evidence on H_e. Then, by composing these solutions (e.g., through a "max" function) the agent would assesses the inferential force of all the favoring evidence on H_e. Because, in this case, there is only one item of favoring evidence, the inferential force is: "Based on the favoring evidence it is very likely that H_e is true."

Through a similar process the agent needs to assess the inferential force of the disfavoring evidence on H_e, and then the likelihood of H_e, based on both the favoring and the disfavoring evidence. In this case, since there is no disfavoring evidence, the assessment of H_e is: "It is very likely that H_e is true." Other elementary hypotheses (e.g. H_d in Fig. 2), are assessed in a similar way, and these assessments are combined, from bottom-up, to obtain the assessment of H_k, as indicated in the right hand side of Fig. 2.

A basic element of the computational theory of evidence-based reasoning is a "substance-blind" ontology of evidence (Schum, 2001a; Schum et al., 2009) which is applicable in every domain. This ontology distinguishes between various types of tangible and testimonial evidence, each with its specific believability credential, as was illustrated above with E_l, an item of demonstrative tangible evidence. As an additional example, the believability credentials of an item of *testimonial evidence based upon direct observation* are the *competence* and the *credibility* of the source. Competence involves *access* (Did this source actually make the observation she claims to have made? Did she have access to the information she reports?) and *understandability* (Did this source understand what was being observed well enough to provide us with an intelligible account?). Credibility involves *veracity* (Is this source telling us about an event she believes to have occurred?), *objectivity* (Did this source base a belief on sensory evidence received during an observation, or did she believe the reported event occurred either because she expected or wished it to occur?), and *observational sensitivity under the conditions of observation* (If the source did base a belief on sensory evidence, how good was this sensory evidence?).

5. INQUIRY-BASED LEARNING

Significant progress has been made in the United States with the development of the *National Science Education Standards* (NRC, 1996), such as the *Science Teaching Standards* that "describe what teachers of science at all grade levels should know and be able to do." These standards call for *inquiry based teaching and learning* which, according to the *National Science Education Standards*: "refers to the diverse ways in which scientists study the natural world and propose explanations based on the evidence derived from their work. Inquiry also refers to the activities of students in which they develop knowledge and understanding of scientific ideas, as well as an understanding of how scientists study the natural world." (NRC, 2000, p. 23). Research indicates that "teaching through inquiry is effective" (NRC, 2000, p. 126) and that the appropriate use of inquiry by the teachers "can have a powerful influence on their students' science learning" (NRC, 2000, p. 128).

And so we illustrate the application of the proposed computational theory of evidence-based reasoning to inquiry-based teaching and learning with an example inspired from (NRC, 2000, pp. 6-11). It is about a natural science class where the students make a surprising observation: "There are two trees outside their classroom window and one has lost its leaves while the other has green leaves" (see Fig. 4). The question is: *What hypothesis would explain this observation?* Through abductive reasoning, the students may generate the following inference chain: "There is evidence that the left tree is dying → It is possible that the left tree is dying → It is possible that too much water at the root of the left tree causes it to die." An alternative hypothesis is that some illness causes the left tree to die.

These two hypotheses will be put to work to collect relevant evidence that will be used to prove or disprove them, as illustrated in the middle of Fig. 4. Here the question is: *Assuming that the hypothesis H_1 is true what other things should be observable?* In this case two other hypotheses are necessarily true: "H_a: There is too much water at the root of the left tree" and "H_b: Too much water at a tree's root causes it to die." Hypothesis H_a directs the students to look for evidence that bears upon it: "Keep the ground around the trees under observation and periodically record the presence of water, noticing that the tree without leaves is almost always standing in water, while the green tree has damp ground and is never standing in water." Similarly, hypothesis H_b directs the students to look for evidence that bears upon it (e.g., a pamphlet from a local nursery entitled "Growing Healthy Plants," and an interview with an expert gardener). As a result, the students find that when plant roots are surrounded by water, they cannot take in air from the space around the roots, and they essentially 'drown.'

A Disciple agent with case studies and examples like the one in Fig. 4 can be used by natural science teachers to teach the processes of inquiry and evidence-based reasoning, including the construction of arguments by assessing the relevance, the believability and the inferential force of evidence. This system could also be used by students to practice inquiry and evidence-based reasoning. For example, a case study could start with some evidence (observation). The students would need to formulate hypotheses that may explain it, use the formulated hypotheses to collect additional evidence (which may include performing various experiments), and determine the most likely hypothesis by evaluating the relevance, believability, and inferential force of their evidence, with the help of the system. Students may compare their analyses, discuss differences, and improve their analyses based on these discussions.

6. SOME COMPLEXITY ISSUES

Evidence-based reasoning is often a very complex task because of a combination of problems that are briefly described below.

Problem 1: There are many dots of different kinds to be connected. As can be seen from Figures 2, 3, and 4, evidence-based reasoning involves connecting many "dots" of different kinds. Some of them are *evidential dots* that concern details in the observable information or data. The second type of dot concerns the *hypotheses* we are trying to prove or disprove, based on our evidence. And the third type of dot, that we call *idea dots*, come in the form of links in chains of reasoning or arguments we construct to link evidential dots to hypotheses. Each of these idea dots refer to sources of uncertainty or doubt we believe to be interposed between our evidence and our hypotheses. This is precisely where imaginative reasoning is involved. The essential task for us is to *imagine* what the evidential dots mean as far as our hypotheses or possible conclusions are concerned. Careful *critical reasoning* is then required to check the logical coherence of sequences of idea dots in our arguments or chains of reasoning. In other words, does the meaning we have attached to sequences of idea dots make logical sense?

Problem 2: Which evidential dots should be connected? Here is where the astonishing complexity of evidence-based reasoning begins to arise because, in order to discover useful hypotheses, we would need to look not just at individual items of evidence, but also at combinations of such items. If we have N evidential dots, then there are $C = 2^N - (N + 1)$ such combinations, which is a very large number in almost any real-world problem.

Problem 3: Which evidential dots should be believed? From some source, a sensor of some sort, or from a person, we obtain an evidential dot saying that a certain event has occurred. But just because this source says that this event occurred does not entail that it did occur. So what is vitally necessary is to distinguish between evidence of an event and the event itself. Complex reasoning is generally necessary in order to assess the believability of an item of evidence.

Problem 4: What are the connections between the evidential dots and the hypotheses? As discussed in the previous sections, all evidence has three major credentials or properties: relevance, believability, and inferential force or weight. No evidence ever comes to us with these three credentials already attached; they must be established by defensible and persuasive arguments linking the evidence to the hypotheses we are considering. Finding such arguments is a very complex problem.

Problem 5: What do our arguments mean? The developed arguments involve the combination of many uncertainties due to the incompleteness, inconclusiveness, ambiguity, dissonance, and varying believability of our evidence. The question remains: How do we assess and combine these assorted uncertainties in complex arguments? The problem is that there are several quite different views among probabilists about what the force or weight of evidence means and how it should be assessed and combined across evidence in arguments. These views include Subjective Bayes, Belief Functions, Baconian, and Fuzzy (Schum, 2001a). Each of these views has something interesting to say, but no one view says it all.

One major objective in developing a computational theory of

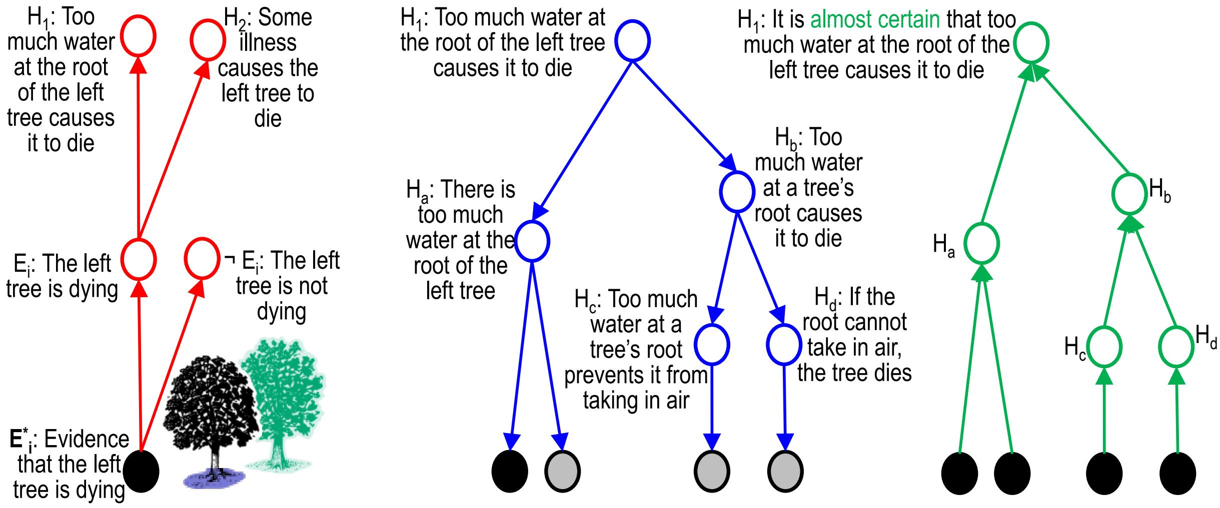

Fig. 4. Inquiry-based learning in a science classroom

evidence-based reasoning is to find solutions to these problems. For instance, research on abductive reasoning may provide solutions to Problem 1 in the form of mixed-initiative strategies for stimulating our imaginative reasoning in hypotheses generation (Schum, 2001b). The hypotheses in search of evidence process, represented by the middle tree in Fig. 2, guides the collection process to identify the dots to be connected in the massive amounts of wide area motion imagery data (Problem 2). The example of believability analysis, illustrated in Fig. 3, shows a partial solution to Problem 3. Argumentation structures based on problem reduction and solution synthesis (see Fig. 3) anchor very clearly the conclusions to the available evidence and provide a partial solution to Problem 4. Further research on combining different types of probability views attempts to find better solutions to Problem 5 (Tecuci et al., 2010b).

7. CONCLUSIONS

We have argued that evidence-based reasoning is at the core of many problem solving and decision making tasks in a wide variety of domains. We have also argued that this is often a very complex task, requiring the development of defensible and persuasive arguments, through imaginative and critical reasoning. These two are justifications for the development of a computational theory of evidence-based reasoning, elements of which have been presented in this paper. This computational theory, which is being developed within the framework of the Scientific Method, is used as a basis for building advanced Disciple cognitive assistants that help students learn critical thinking skills for evidence-based reasoning, through a hands-on approach (Tecuci et al., 2010b) and support users in coping with the complexity of evidence-based reasoning in a variety of domains.

We think that evidence-based reasoning should be studied as a general critical thinking skill, at all levels of education, from the elementary school to the university, and we propose Disciple-based agents as supporting educational tools.

ACKNOWLEDGEMENTS

This research was performed in the Learning Agents Center and was partially supported by the National Geospatial-Intelligence Agency and by the Department of Defense. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and opinions expressed in this article are those of the authors and do not necessarily reflect the official policy or position of any agency of the U.S. Government.

REFERENCES

Betham, J. (1810). *An Introductory View of the Rationale of the Law of Evidence for Use by Non-lawyers as well as Lawyers* (vi *works* 1-187 (Bowring edition, 1837-43) originally edited by James Mill circa 1810).

Murphy, P. (2003). *Evidence, Proof, and Facts: A Book of Sources.* Oxford University Press, Oxford, UK.

Nilsson, N.J. (1971). *Problem Solving Methods in Artificial Intelligence.* NY: McGraw-Hill.

Noon, R.K. (2009). *Scientific Method: Applications in Fact Investigation and Forensic Science*. CRC Press, Taylor & Francis, Boca Raton, Florida.

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.

Peirce C.S. (1898). *Reasoning and the Logic of Things*, 1898. In Ketner K., ed., Harvard University Press, Cambridge, MA, 1992.

Peirce C.S. (1901). Abduction and Induction, 1901. In *Philosophical Writings of Peirce*, Buchler J., ed., Dover, New York, NY, 1955, pp. 150-156.

Schum, D. (1987). *Evidence and Inference for the Intelligence Analyst* (Two Volumes). Lanham, MD: University Press of America.

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

Schum, D.A. (2001b). Species of Abductive Reasoning in Fact Investigation in Law. *Cardozo Law Review*, 22 (5-6), pp. 1645–1681.

Schum, D.A. (2009). Science of Evidence: Contributions from Law and Probability. *Law Probab Risk* **8** 197–231.

Schum, D., Tecuci, G., Boicu, M., Marcu, D. (2009). Substance-Blind Classification of Evidence for Intelligence Analysis. *Proc. Conf. "Ontology for the Intelligence Community,"* Fairfax, VA, 20-22 October.

Tecuci, G. (1988). *DISCIPLE: A Theory, Methodology and System for Learning Expert Knowledge*. Thèse de Docteur en Science, University of Paris-South, France.

Tecuci, G. (1998). *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies*, San Diego, CA: Academic Press.

Tecuci, G., Boicu, M., Cox, M.T. (2007). Seven Aspects of Mixed-initiative Reasoning: An Introduction to the Special Issue on Mixed-initiative Assistants. *AI Magazine* **28(2),** 11–18.

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 (JICRD)*, July.

Tecuci, G., Schum, D.A., Boicu, M., Marcu, D., Hamilton, B. (2010a). Intelligence Analysis as Agent-Assisted Discovery of Evidence, Hypotheses and Arguments. In: Phillips-Wren, G., Jain, L.C., Nakamatsu, K., Howlett, R.J. (eds.) *Advances in Intelligent Decision Technologies*, pp.1-10. Springer-Verlag Berlin Heidelberg.

Tecuci, G., Schum D.A., Boicu, M., Marcu D. (2010b), *Introduction to Intelligence Analysis: A Hands-on Approach with TIACRITIS*, 220 pg, George Mason Univ.

Toulmin, S. E. (1963). *The Uses of Argument*. Cambridge University Press.

Twining, W. (2003). Evidence as a Multi-Disciplinary Subject. Law, Probability and Risk, Vol.2, No.2, 91-105.

Wigmore J.H. (1913). The Problem of Proof. *Illinois Law Review*, Vol. 8, No. 2, 1913, pp. 77 – 103.

Wigmore, J.H. (1937). *The Science of Judicial Proof*. Boston, MA: Little, Brown & Co.