Accepted manuscript of the paper published in
Intelligent Decision Technologies, Volume 8, Issue 3, July 2014, Pages 231-250.

Toward Cognitive Assistants for Complex Decision Making Under Uncertainty

D.A. Schum, G. Tecuci, D. Marcu, M. Boicu
Learning Agents Center, George Mason University, Fairfax, VA, USA

Abstract. Discussed in this paper is a quite unique and novel intelligence decision technology resting upon three systems we have called Disciple-LTA [Learning, Teaching and Assistance], TIACRITIS [Training Intelligence Analysts Critical Reasoning Skills], and Disciple-CD [Connecting the Dots]. We have so far applied these systems to complex intelligence inferences based on masses of evidence of many different kinds and coming from many different sources. This paper discusses the extension of these systems to be valuable decision support assistants that are capable of helping analysts to answer the two fundamental questions regarding decisions made in the face of uncertainty: what's at stake?, and what are the odds? The stakes question concerns the value or utility of decision consequences and the odds question concerns the probability of these possible consequences. We discuss the requisite ingredients of defensible and persuasive decisions and problems associated with the discovery of these ingredients in a world that keeps changing all the time. But we also consider the constraints facing intelligence analysts who so often have limited time for decisions and who also have deficiencies regarding the availability of information supporting requisite value and probability judgments. Conventional approaches to decision analysis are usually not helpful in the face of these constraints. We offer simplified methods for assessing both value and probability judgments and a simplified method for combining these judgments in the selection of a course of action that does take account of the requisites for defensible and persuasive decision and analysis. In the process, we illustrate our methods with a very complex analysis involving the possible proliferation of nuclear weapons.

Keywords: decision making under uncertainty, decision making under time and information constraints, intelligence analysis, decision rules, cognitive assistants, discovery, combining judgments, proliferation of nuclear weapons.

1 Introduction

For many years we have worked on creating a theory, methodology, and tools for the development of knowledge-based cognitive assistants that:

- learn complex problem solving expertise directly from subject matter experts,
- support experts and non-experts in problem solving,
- teach their problem solving expertise to students.

This approach relies on developing powerful learning agents that can be taught by the subject matter experts (who do not have computer science or knowledge engineering experience) in ways that are similar to how the experts would teach students or apprentices, by explaining problem solving examples to them, and by supervising and correcting their problem solving behavior. Because such agents learn to replicate the problem-solving behavior of their human experts, we have called them Disciple agents (Tecuci, 1988; 1998; Boicu, 2002; Marcu, 2009; Tecuci et al., 2002, 2005; 2008a, 2013).

The long term goal of the Disciple approach is to contribute to a new revolution in the use of computers by enabling typical computer users to develop their own cognitive assistants. Thus, non-computer scientists will no longer be only users of generic programs developed by others (such as word processors or Internet browsers), as they are today, but also agent developers themselves. They will be able to train their personal Disciple assistants to help them with their increasingly complex tasks in the knowledge society, which should have a significant beneficial impact on their work and life.

Advancing toward this goal, we have more recently researched a Computational Theory of Intelligence Analysis (Tecuci et al., 2010a; 2011), and we have significantly expanded the Disciple approach to incorporate elements of this theory, enabling the rapid development of knowledge-based cognitive assistants for intelligence analysis and other evidence-based reasoning domains (Boicu et al., 2011). This has led to the development of Disciple-LTA (Tecuci et al., 2008b), TIACRITIS (Tecuci et al., 2010b), and Disciple-CD (Tecuci et al., 2013), which are cognitive assistants for intelligence analysis. All these systems have a large amount of knowledge about the properties, uses, discovery, and marshaling of the evidence, as well as domain-specific analytic expertise learned from expert analysts, which they apply to the current analysis process. TIACRITIS is an improvement of the intelligence analysis functionality of Disciple-LTA, but relies on the knowledge engineering and learning modules of Disciple-LTA to develop its knowledge base. Disciple-CD is a further improvement of TIACRITIS. As a result, Disciple-CD significantly supports intelligence analysts in coping with the astonishing complexity of generating and analyzing hypotheses about events in the world, to inform the decision-making process.

With the development and partial implementation of the Computational Theory of Intelligence Analysis, a natural next step is to extend the Disciple approach to enable the development of cognitive assistants for complex decision making under uncertainty. Indeed, as discussed in this paper, decisions made in the face of uncertainty always involve two basic questions:

- *What is at stake? or What are the possible consequences of our alternative actions or decisions, and the utility or value we place on them?*
- *What are the odds? or What is the probability of various states of the world that can produce different patterns of consequences depending on what decision or action we select?*

The Computational Theory of Intelligence Analysis addresses the second question. The purpose of this paper is to also address the first question so that we may consider the entire decision process.

Providing productive and acceptable assistance to *decision makers* in any context is not easy, as anyone who attempts to do so quickly discovers. Indeed, decisions made in the face of uncertainty have many probability and value-related ingredients that require very complex judgments; there is always concern about how these judgments might best be made and how they should be combined in order to select a course of action. These matters have been the subject of careful study for many years by persons throughout academia. There is extended academic discussion, and much controversy, about what constitutes "optimal" or "rational" methods for probabilistic and value-related assessments and their aggregation.

This paper describes our attempt to design and develop a decision-support system having what we believe are truly unique and very valuable characteristics. In describing our approach to this system, we first focus on the users of our system, whether they are decision makers themselves or persons providing direct assistance to decision makers. From years of experience we know that the users of our system, whoever they might be, have some constraints or limitations they must face every day. The first common constraint is temporal in nature. Decision makers often face severe limitations on the amount of time they have to make choices, even in crucial decision situations which involve the lives and resources of many people. The second constraint involves limitations on the information the decision makers need to support the probability or value-related judgments they must make, judgments which often require various forms of ancillary information that might not be readily available or are very difficult and time-consuming to obtain.

The Disciple-based approach that we will describe allows the user to decompose both probability and value-related judgments according to a common strategy called "divide and conquer." By itself, this is hardly a novel approach since virtually every decision aid employs this approach. But Disciple allows the user to decompose (or "drill down") a probability or value-related task to various levels, depending on the temporal or resource limitations the user faces. In some cases, such as when the formal requirements are quite severe but judgments have to be made quickly, this drilling down will be very "shallow." We have always been concerned about whether these shallow drilling methods allow judgments that are defensible in light of user constraints and the requisites of defensible and persuasive choice processes.

Another important matter is the fact that the world is continually changing so that we need to continually revise our intended inferences and choices. This reflects a concern we have always had in our research with Disciple-LTA, TIACRITIS, and Disciple-CD. Our systems allow users to accommodate to these non-stationarities in various ways.

One very good way of assessing which stakes issues and attributes we ought to consider involves the construction of examples of situations in which we have non-trivial choices to make having serious consequences, and where these choices are made in the face of very complex states of the world, that keep changing all the time. In this paper we use a real life decision situation that offers as great a level of complexity as any we might imagine. In particular, we use as examples factors involved in our leaders' possible choices concerning the apparent threats posed by the leadership of country "Purple" who, apparently, is attempting to develop nuclear weapons.

First, in Section 2, we provide an overview of decision-making under uncertainty, to introduce the ingredients of the decision-making process that need to be discovered through imaginative reasoning. Then, in the following sections, we will illustrate each step of this process, and the corresponding methods used in the Disciple approach. We conclude the paper with a summary of the main contributions of the proposed approach.

2 Decision-making Under Uncertainty

As we noted, decisions made in the face of uncertainty always involve two basic questions, *"What is at stake?"* and *"What are the odds?"* The *stakes question* addresses the possible consequences of actions or decisions we are considering and the utility or value we place on them. Different decisions or actions will of course imply different patterns of consequences, the utility or value of each of these consequences we will need to assess. The *odds question* involves our assessing the probability of various states of the world that can produce different patterns of consequences depending on what decision or action we select.

We first identify several ingredients of the decision-making process that need to be discovered through imaginative reasoning. In so many works it is assumed that all these ingredients are already in existence, as if they had been supplied by someone to the person(s) facing a decision task.

First the decision-maker has to state the goals or objectives to be achieved: $G_1, G_2, \ldots, G_i, \ldots, G_s$

Our decision maker now has to generate or imagine what different options, actions, or choices should be considered in order to achieve these goals: $O_1, O_2, \ldots, O_i, \ldots, O_n$.

The decision maker also has to imagine possible conditions or world states that may give rise to different consequences for any action that might be taken. Usually, these possible world states are referred to as *hypotheses*, $H_1, H_2, \ldots, H_j, \ldots, H_t$. Here is where uncertainty enters our picture because each hypothesis H_j will have a certain probability P_j to occur.

The next task concerns identifying the consequences of the options being considered. These consequences emerge by the joint consideration of an option and a world state (hypothesis): If option O_i is chosen and the world is in state H_j, then the outcome or result will be the consequence $C_{i,j}$.

Now we want to have measurable indications of the extent to which our consequences $C_{i,j}$ meet or fail to meet our goals G_r. For this we have to identify a list of consequence attributes $A_1, A_2, \ldots, A_m, \ldots, A_r$. Each *attribute* A_m is a *measurable* indication of the extent to which a consequence is in accordance with one of our stated objectives or goals.

We make the following quite useful notational distinction. Suppose we have the consequence attributes $A_1, A_2, \ldots, A_r$. Considering some particular consequence $C_{i,j}$, we have to measure to see how much of each attribute it has. We list the levels or states of these attributes for this consequence as: $A^*_1(C_{i,j})$, $A^*_2(C_{i,j})$, …, $A^*_r(C_{i,j})$. Thus, for every consequence $C_{i,j}$, we have r measured values which, for simplicity, we denote as $C^*_{i,j}$: A^*_1, A^*_2, …, A^*_m, …, A^*_r.

Next comes an important value judgment. Considering the just measured attribute A^*_m as an example, we must judge the value $V(A^*_m)$ of having this amount of A_m. We need such judgments for each of our r measured attributes of $C^*_{i,j}$.

Supposing that we now have a value assessment $V(A^*_m)$ for each attribute of consequence $C^*_{i,j}$, one other judgment is necessary. The decision maker needs to assess the relative importance of each of the measured attributes A^*_1, A^*_2, …, A^*_m, …, A^*_r. So, what we have will be a vector of weighted value judgments:

$$W_1V(A^*_1), W_2V(A^*_2), \ldots, W_mV(A^*_m), \ldots, W_rV(A^*_r)$$

Next we have to combine these weighted value judgments to obtain the composite (or total) *utility* or *value* $V(C^*_{i,j})$. This will tell us the extent to which we will achieve, or fail to achieve, our goals if we choose option O_i when the world is in state H_j.

Table 1 summarizes the ingredients necessary to determine the value of an option or action. We have vector consequences associated with option O_i, as given by the $C^*_{i,j}$ values. But we have t values of $C^*_{i,j}$, one under each of the t possible relevant states of the world H_j. We also have composite value determinations $V(C^*_{i,j})$, for every vector consequence $C^*_{i,j}$ and the probabilistic ingredients, shown by the vector of P_j values. What these probabilities mean is that, in answering the question "What are the odds?" we have combined a mass of evidence in order to determine the probability of each state of the world we are considering. Therefore $P_j = P(H_j|E^*)$, where E* is all the evidence we have considered.

Table 1. The ingredients necessary to determine the value of an option or action.

	P_1	P_2	…	P_j	…	P_k	…	P_t
	H_1	H_2	…	H_j	…	H_k	…	H_t
O_i:	$C^*_{i,1}$	$C^*_{i,2}$	…	$C^*_{i,j}$	…	$C^*_{i,k}$	…	$C^*_{i,t}$
	$V(C^*_{i,1})$	$V(C^*_{i,2})$	…	$V(C^*_{i,j})$	…	$V(C^*_{i,k})$	…	$V(C^*_{i,t})$

What we have to do next is to combine all these ingredients in order to determine $V(O_i)$. And finally, we have to choose the best option.

Up to this point we have shown the main steps that need to be performed in order to make decisions under uncertainty. What we have not shown is how to perform these steps. This involves some very difficult judgments.

To facilitate the presentation of the proposed methods, we use as examples factors involved in our leaders' possible choices concerning the threats posed by the leadership in country Purple, in their apparent pursuit of nuclear weapons. Our examples will illustrate situations in which we consider what Purple might do and what actions we might take either to prevent them from developing nuclear weapons or to make them regret having developed such weapons.

2.1 Goals or Objectives

We will start by considering a major source of controversy in decisions which involve multiple stakeholders. In many situations, the goals or objectives to be served, and their relative importance, are never considered carefully. When they are, it often happens that important stakeholder groups are not consulted regarding their views on these objectives and their relative importance.

It would be difficult to find a decision having more potential stakeholders than decisions having international consequences of responses to perceived nuclear threats. The resulting confrontations are known to spread rapidly and end up involving many countries around the world.

When the interests of many stakeholder groups are considered, we expect to find many different patterns of goals or objectives that should be pursued. What we also expect is that many of these goals are *conflicting*; we can't have both A and B at the same time. To get the amount of A we want, we may have to give up some of the B we also want.

A related problem here concerns assessing the relative importance of these goals or objectives. Stakeholders may well agree about what goals or objectives should be pursued in a decision but may disagree, often substantially and heatedly, about their relative importance.

Additionally, the goals or objectives are often stated so vaguely that we will seldom be able to tell how well our choices have achieved them. In various ways, our initial goals or objectives need to be parsed carefully to make them as specific as possible.

There is yet another problem to consider, the decision maker's goals or objectives can be *non-stationary*. Peoples' objectives can change over time.

Finally, and this is very important, are our asserted goals or objectives really ones that are in our best interests and that we should be pursuing?

Thus, one research problem is how a Disciple cognitive assistant can help a decision maker in the process of asserting goals or objectives in light of these issues.

The following is an initial listing of goals or objectives our leaders might consider in decisions concerning possible nuclear weapons threats by Purple. We have also listed our own *ordinal* assessment of their relative importance; i.e., we have ranked them from **1** (most important) to **10** (least important).

- **G_1:** Preserve human lives. **1 (tied)**
- **G_2:** Promote stable and peaceful relations among all countries in the region of Purple. **4**
- **G_3:** Don't make the lives of peaceful and productive Purple civilians any worse. **7**
- **G_4:** Stop the proliferation of nuclear weapons. **3**
- **G_5:** Protect the safety and interests of our military and civilian personnel. **1 (tied)**
- **G_6:** Take actions that are within our present economic limitations. **8**
- **G_7:** Preserve our country's leadership role in international affairs. **9**
- **G_8:** Take actions that are feasible responses by our military. **5**
- **G_9:** Take actions that are politically acceptable to our citizens. **6**
- **G_{10}:** Preserve the flow of goods from countries in the region of Purple. **10**

2.2 Options, Actions, or Choices

Our decision maker now has to generate or imagine what different options, actions, or choices should be considered. This is an important imaginative reasoning task since, unless the decision is one that recurs, new options must be considered. However, there is an important attribute about generated options, namely that they must be *admissible*. Admissibility means that an option being considered corresponds with one or more of our stated goals or objectives. We also note that these options/actions are contingent and may need to be revised, eliminated, or enhanced as we learn new things. "Now" in these options refers to some unspecified future time.

O_1: Defer a Decision. At this "now" time it might make perfect sense to decide not to make any specific decision but to *buy more information*. Deciding to postpone a decision is itself a decision that may result in a better decision later when a specific choice is required. Of course there are risks associated with postponing specific decisions.

O_2: Suspend Military Plans and Focus on Non-Confrontational Measures. This action amounts to increasing our dialogue with Purple's leaders in an effort to remove their aspirations to develop nuclear weapons. We might call this the "peaceful action" which we could accompany with not pressing for any additional sanctions against Purple.

O_3: Finalize Plans for Military Action against Purple to be Taken in the Near Future. This action calls for a suspension of attempts to get the Purples to quit developing an arsenal of nuclear weapons. If we decide to attack Purple's suspected nuclear weapons development facilities, who will be with us? We will suppose here that we will not take unilateral military action against Purple.

2.3 Hypotheses or Possible World States

Here is where uncertainty enters our picture and we need to answer the question: What are the odds? The decision maker has to imagine possible conditions or world states that may give rise to different consequences for any action that might be taken. Usually, these possible world states are referred to as *hypotheses*.

As mentioned in Section 1, we have already made significant progress on developing and implementing a general computational approach for answering the "odds" question with a Disciple cognitive assistant. Here we only provide a summary illustration of our approach which is presented in more detail in various papers (Tecuci et al., 2013).

Within the framework of the scientific method, we view intelligence analysis as ceaseless discovery of evidence, hypotheses, and arguments in a non-stationary world, involving collaborative processes of evidence in search of hypotheses, hypotheses in search of evidence, and evidentiary testing of hypotheses (see Figure 1).

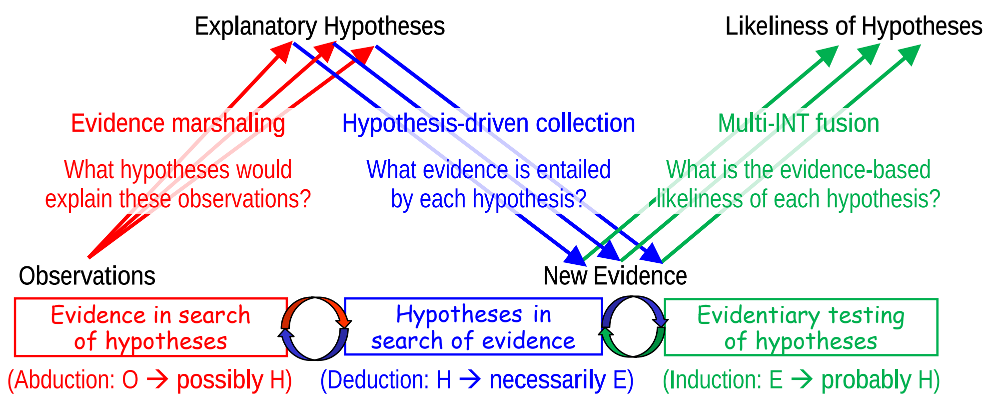

Figure 1. Intelligence analysis as discovery of evidence, hypotheses and arguments.

Through *abductive* (*imaginative*) reasoning (which shows that something is *possibly* true) the analyst and Disciple-CD generate alternative hypotheses that explain their observations (see the left side of Figure 1). Through *deductive* reasoning (which shows that something is *necessarily* true) they use these hypotheses to generate new lines of inquiry and discover new evidence (see the middle of Figure 1). And through *inductive* reasoning (which shows that something is *probably* true) they test each of these hypotheses with the discovered evidence and select the most likely one (see the right side of Figure 1). As new evidence is discovered, new hypotheses are being generated or old ones are being updated which, in turn, drive new processes of evidence discovery. Similarly, incomplete or inconclusive assessments of the alternative hypotheses require searching for additional evidence. These feedback loops are represented at the bottom of Figure 1.

Consider, for example, the discovery of evidence that Purple is producing enriched uranium. The question is: *What hypotheses would explain these observations?* One hypothesis is that Purple is developing nuclear weapons. But this is a very general hypothesis that needs to be refined by considering various matters, such as: kind of weapons (e.g. nuclear or thermonuclear), yield, delivery methods, mission (e.g., tactical or strategic), number, what weapons can Purple afford to construct, when will they be ready to use, when will Purple attempt to use them. Many world states or hypotheses result from considering various *possible combinations* of such alternatives. And, of course, there is the additional hypothesis that Purple is not, in fact, developing any nuclear weapon. In this paper, however, we will only consider the following hypotheses which are obviously not exhaustive and might not even be mutually exclusive.

H_1: Purple is not "now" developing any nuclear weapons and has no intention of developing any in the foreseeable future.

H_2: Purple began developing nuclear weapons in the kiloton yield range but has "now" suspended these developments, at least temporarily.

H_3: Purple is "now" in the act of developing several nuclear weapons having a yield of no more than 20 kilotons each and, within a year's time, will be able to missile-launch these weapons at targets within a 1500-mile range.

H_4: Purple is "now" in the act of developing several nuclear weapons having yields of at most 40 kilotons each and, within two years' time, will be able to missile-launch these weapons at targets within a 2500-mile range.

The first two hypotheses represent possibilities saying that Purple does not "now" present any nuclear weapons threats. Hypothesis H_1 says they never presented any nuclear threat; hypothesis H_2 says that the nuclear threat Purple posed has "now" disappeared but may emerge again in the future. Hypotheses H_3 and H_4 are possible threat levels based upon our guesses about numbers of weapons, their yields, development times, and missile ranges.

Having generated alternative hypotheses (left side of Figure 1), the next step is to search for evidence to assess them (middle of Figure 1). The basic idea is putting the generated hypotheses at work to guide the collection of relevant evidence. The question is: *"Assuming that the hypothesis is true, what other events or entities should be observable?"* The deductive reasoning process for answering this question successively reduces the assessment of the top-level hypothesis to the assessments of simpler hypotheses, ultimately resulting in precise evidence collection guidance, as illustrated in Figure 2 with the hypothesis H_3. Notice that H_3 is first reduced to 6 simpler hypotheses:

- **H_{31}**: Purple is "now" in the act of developing nuclear weapons.
- **H_{32}**: Purple is developing several nuclear weapons.
- **H_{33}**: Purple's nuclear weapons have yields no greater than 20 kilotons.

- **H$_{34}$**: Purple's nuclear weapons will be ready in a year's time.
- **H$_{35}$**: Purple's nuclear weapons will be missile-launched.
- **H$_{36}$**: Purple's nuclear weapons will be launched at targets within a 1500 mile range.

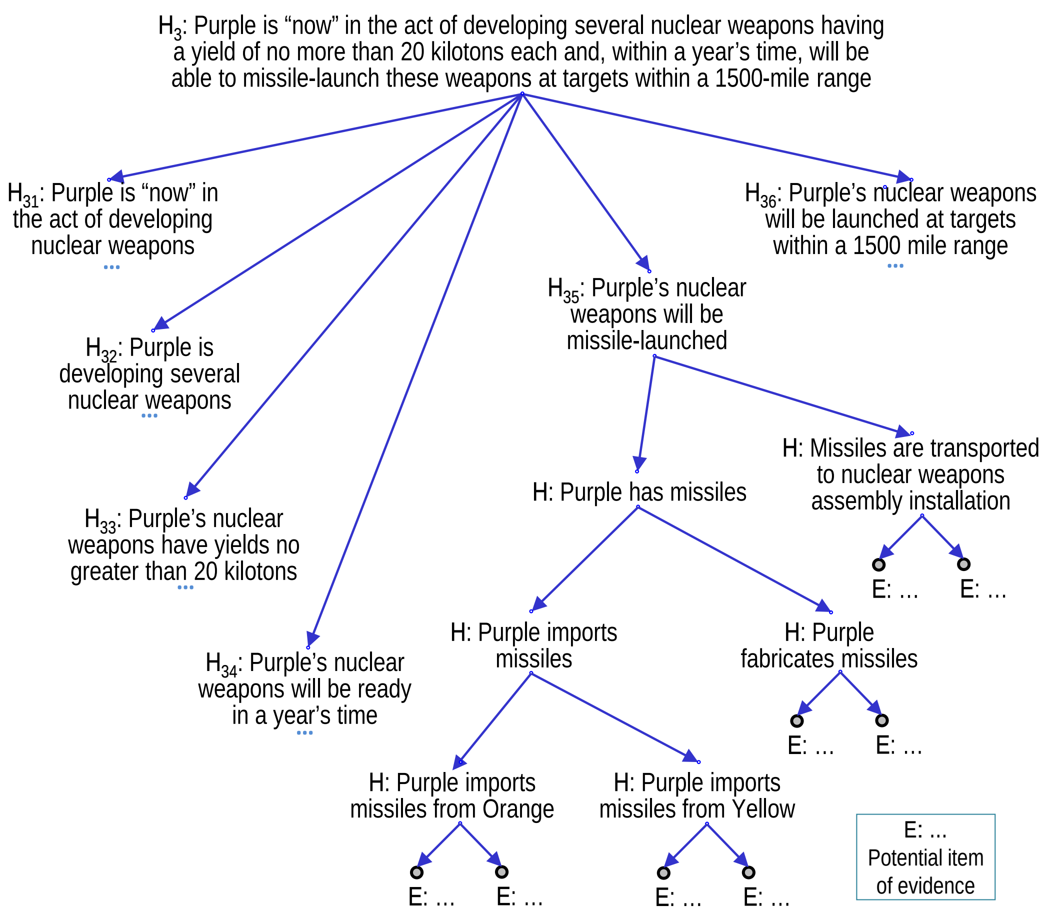

Figure 2. Hypotheses in search of evidence.

Each of these hypotheses is further reduced to simpler and simpler hypotheses, as illustrated with the hypothesis "**H$_{35}$**: Purple's nuclear weapons will be missile-launched." These simpler hypotheses guide us in the collection of evidence (shown as shaded circles in Figure 2) that may prove or disprove them. For example, the hypothesis "Purple imports missiles" guides us to collect evidence that Purple may be importing them from Orange or, perhaps, Yellow. Similarly, the hypothesis "Purple fabricates missiles" guides us to look for evidence of corresponding Purple's installations. We are also guided to look for evidence that "Missiles are transported to the nuclear weapons assembly installation". We may find MOVINT evidence of truck traffic between a seaport in Purple (where the imported missiles could be unloaded) and the Purple's nuclear weapons assembly installation. Or we may discover evidence of a Purple installation that fabricates missiles, as well as MOVINT evidence of truck traffic between this installation and the nuclear weapons assembly installation.

Having discovered evidence relevant to the alternative hypotheses, the next question is: *"What is the evidence-based likeliness of each hypothesis?"* (see the right side of Figure 1). Here Disciple-CD uses an approach grounded in the problem reduction representations developed in the field of Artificial Intelligence (Nilsson, 1971; Powell and Schmidt, 1988; 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, 2001). In this approach, which is illustrated in Figure 3:

- The problem of assessing a complex hypothesis is successively reduced to the assessment of simpler and simpler hypotheses (as also illustrated in Figure 2).
- The leaf hypotheses are assessed based on the available evidence.
- Finally, these assessments are successively combined, from bottom-up, to obtain the assessment of the top level hypothesis.

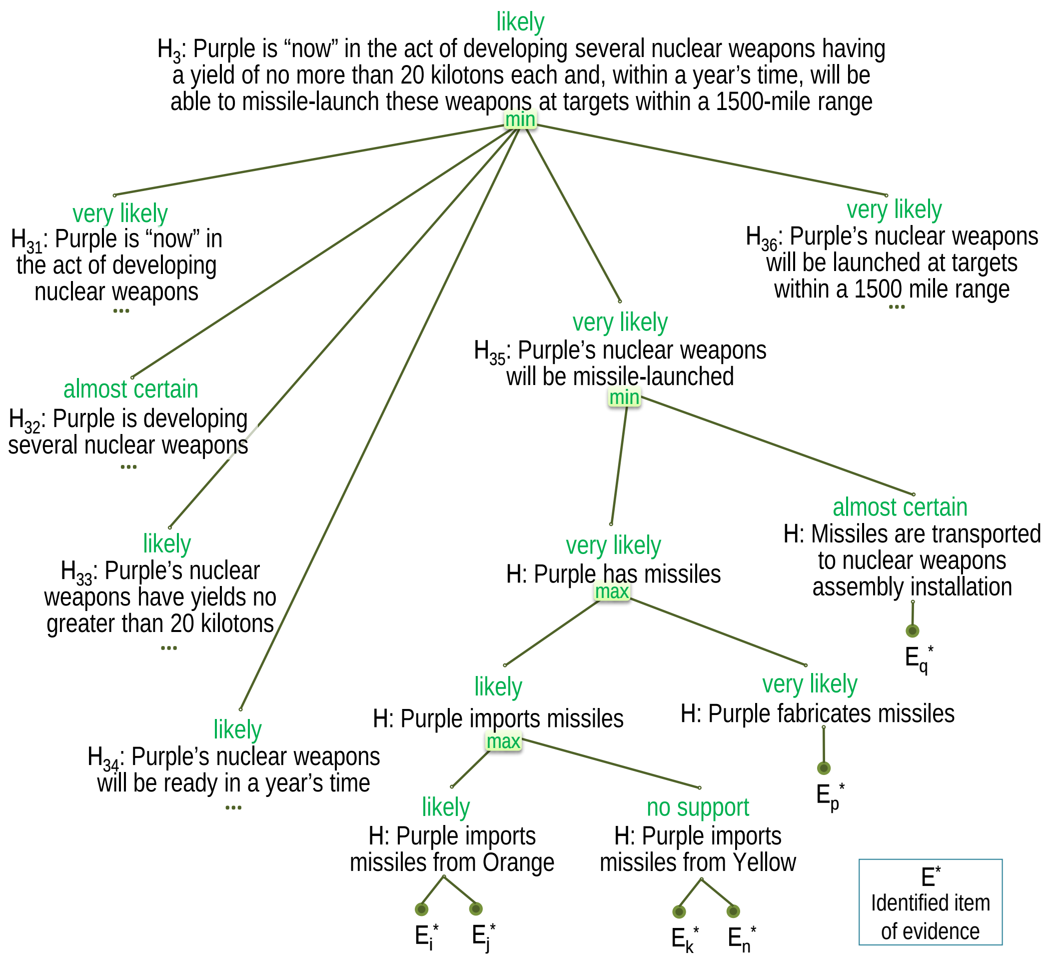

Figure 3. Evidence-based hypotheses analysis through reduction and synthesis.

A conventional ways of assessing and combining probabilities in evidential reasoning tasks is to use Bayes' rule (Pearl, 2009). But the use of Bayes' rule in the complex argument structures comes with a cost that many analysts will not be willing to pay. There will be enormous numbers of probabilities to assess and many of them might be difficult or impossible to evaluate. Add to this the fact that in most inference about matters of national security, the events of concern are singular, unique, or one-of-a-kind. Consequently, there will be no statistics to support probability assessments. Consequently, Disciple-CD employs a much simpler and easier to use approach to assessing and combining probabilities, which is based on ideas from the Baconian and Fuzzy systems of probabilistic reasoning. The Baconian system we employ was developed by Professor L. Jonathan Cohen, late of Queens College, Oxford (Cohen, 1977). This probability system rests on ideas generated centuries ago by Sir Francis Bacon. This Baconian system is especially attractive since it is the only probability system known to us that takes explicit account of how complete is our evidence. We must consider how many relevant questions remain unanswered by our existing evidence. As in the Fuzzy system (Zadeh, 1983), the probabilities are expressed in words rather than in numbers. In particular, we use the ordered symbolic probability scale from Table 2 (but more or less refined scales could also be defined). As in the Baconian system (Cohen, 1977), no support for a hypothesis means that we have no basis to consider that the hypothesis might be true. However, we may later find evidence that may make us believe that the hypothesis is very likely, for instance.

Table 2. Ordered symbolic probability scale.

no support < likely < very likely < almost certain < certain

To assess the hypotheses, we first need to attach each item of evidence to the hypothesis to which it is relevant. Then we have to establish the *relevance*, the *believability*, and the *inferential force or weight* of evidence on the corresponding leaf hypotheses. Next we have to combine the probabilistic judgments of the leaf hypotheses into the probabilistic judgments of the upper-level hypotheses and the top-level hypotheses, as shown in Figure 3. As can be seen in this example, Disciple-CD uses simple min/max combination rules, consistent with both the Fuzzy and the Baconian systems of probability. What combination function is used depends on whether the sub-hypotheses of a hypothesis represent one or several sufficient conditions, or just indicators.

Disciple-CD has a large amount of knowledge about the properties, uses, discovery, and marshaling of the evidence, as well as domain-specific analytic expertise learned from expert analysts, which help its user to perform better analyses faster. In particular, the analysis of other hypotheses (e.g., H_4) will contain sub-hypotheses similar to those in Figure 3 (e.g., H_{31}, H_{32}, and H_{35}). In such cases Disciple-CD will generate the corresponding argumentation structures by using rules learned from the analysis of H_4.

Additionally, if new evidence is discovered, it can be easily added in the argumentation structures, and Disciple-CD will automatically update the assessments of the hypotheses.

Finally, we add here that the use of the Baconian/Fuzzy methods have a direct impact on the methods we use to answer the "stakes" question, as will be discussed in the following.

2.4 Generation of Consequences

The next task involving discovery and imaginative reasoning concerns identifying the consequences of the options being considered. These consequences emerge by the joint consideration of an option and a world state (hypothesis). Given any action O_i and hypothesis H_j, $C_{i,j}$ is the consequence of taking action O_i if the world is in state H_j. Given the four hypothesized states of the world and the three options we have generated, there are "now" twelve consequences (C) to be considered (see Table 3).

Table 3. Options, hypotheses and consequences in the Purple example.

	H_1	H_2	H_3	H_4
O_1:	$C_{1,1}$	$C_{1,2}$	$C_{1,3}$	$C_{1,4}$
O_2:	$C_{2,1}$	$C_{2,2}$	$C_{2,3}$	$C_{2.4}$
O_3:	$C_{3,1}$	$C_{3,2}$	$C_{3,3}$	$C_{3,4}$

Now, clearly these twelve consequences have many attributes that we will have to consider. But first, we can give the following brief verbal account of these consequences that will help us as we think about their attributes and the utility or value we associate with them:

$C_{1,1}$: Here our decision is only to buy more information about Purple, when it is true that Purple is not "now" developing any nuclear weapons and has no intentions of developing any in the foreseeable future. But perhaps its present intentions will change.

$C_{1,2}$: Here our decision is only to buy more information about Purple, when it is true that Purple has been developing nuclear weapons in the kiloton yield range but has "now" suspended these developments at least temporarily. The information we are buying will help us decide how temporary their cessation of weapons development will be.

$C_{1,3}$: Here our decision is only to buy more information about Purple. But it happens that Purple is "now" in the act of developing several nuclear weapons having a yield of no more than 20 kilotons each and, within a year's time, will be able to missile-launch these weapons at targets within a 1500-mile range. In this case we could not defer other decisions for too long a time if Purple will have a nuclear weapons arsenal within a year's time.

$C_{1,4}$: Here our decision is only to buy more information about Purple. But it happens that Purple is "now" in the act of developing several nuclear weapons having yields of at most 40 kilotons and, within two year's time, will be able to missile-launch these weapons at targets within a

2500-mile range. Here we may have a slightly longer time to defer making other decisions, but making other decisions may be more urgent given the advanced Purple capabilities.

$C_{2,1}$: In this case we decide to suspend plans for military actions against Purple and focus on non-confrontational measures, when it is true that Purple is not "now" developing any nuclear weapons and have no intention of developing any in the foreseeable future. In this case we would be relying on Purple's leadership to continue to pursue rational and peaceful uses for nuclear materials.

$C_{2,2}$: In this case we decide to suspend plans for military actions against Purple and focus on non-confrontational measures. However, it is true that Purple has been developing nuclear weapons in the kiloton yield range but has "now" suspended these developments at least temporarily. This decision would rest on faith that our non-confrontational measures would convince Purple to suspend their development of nuclear weapons permanently.

$C_{2,3}$: In this case we decide to suspend plans for military actions against Purple and focus on non-confrontational measures. But it happens that Purple is "now" in the act of developing several nuclear weapons having a yield of no more than 20 kilotons each and, within a year's time, will be able to missile launch these weapons at targets within a 1500-mile range. If we discovered these developments, we would have to consider revising this decision.

$C_{2.4}$: In this case we decide to suspend plans for military actions against Purple and focus on non-confrontational measures. But it happens that Purple is "now" in the act of developing several nuclear weapons having yields of at most 40 kilotons and, within two year's time, will be able to missile-launch these weapons at targets within a 2500-mile range. If we discovered these developments, we would have a bit more time to consider revising this decision.

$C_{3,1}$: In this case we finalize plans for military action to be taken against Purple in the near future. But it happens that Purple is not "now" developing any nuclear weapons and has no intention of developing any in the foreseeable future. In this case, we seem to have been misled about Purple's nuclear weapons aspirations and capabilities. If our plans for military action leaked out, we would suffer widespread condemnation for war-mongering.

$C_{3,2}$: In this case we finalize plans for military action to be taken against Purple in the near future. But it happens to be true that Purple has been developing nuclear weapons in the kiloton yield range but have "now" suspended these developments at least temporarily. In this case, we might consider placing our military plans on hold.

$C_{3,3}$: In this case we finalize plans for military action to be taken against Purple in the near future. It happens that Purple is "now" in the act of developing several nuclear weapons having a yield of no more than 20 kilotons each and, within a year's time, will be able to missile-launch these weapons at targets within a 1500-mile range. If we knew these things to be true, we could be ready to engage Purple militarily on fairly short notice if this seemed advisable.

$C_{3,4}$: In this case we finalize plans for military action to be taken against Purple in the near future. But in this case it happens that Purple is "now" in the act of developing several nuclear weapons having yields of at most 40 kilotons and, within two year's time, will be able to missile-launch these weapons at targets within a 2500-mile range. If we knew these things to be true, we might have to revise our military plans, but we would not have to start from the beginning.

The basic trouble here is that we don't know for sure which one of our four hypotheses is "now" true. To be fair, we do have what we regard as relevant evidence on these hypotheses and Disciple-CD has allowed us to assess the believability and force of the evidence we "now" have. Suppose we could be certain that one of these hypotheses, say H_3, was "now" true. This would simplify our decision task, but many difficulties would still remain. If H_3 were true, we would have consequence $C_{1,3}$ if we chose option O_1; consequence $C_{2,3}$ if we chose option O_2; and consequence C_{33} if we chose option O_3. All we would have to do here is determine which of these three consequences we liked the best and then select the option that produces this consequence. For example, if we liked $C_{2,3}$ best, then we would select option/action O_2.

Here is problem #1: *These consequences have many attributes.* Deciding which one of these three consequences we prefer, may not be so easy.

Now, here is problem #2: We will not know for sure which hypothesis is true. *This makes it necessary for us to consider how much we like or dislike each one of all twelve of our consequences.* To do this, we must try to define their *essential attributes*.

2.5 Consequence Attributes

Where do these consequence attributes come from? They come from careful consideration of our decision goals or objectives. There are two ways of generating consequence attributes, a *top-down* approach and a *bottom-up* one.

The *top-down approach* resembles our "hypotheses in search of evidence" where we attempt to infer observable evidence relevant to these hypotheses. As we have seen, this is a hierarchical inference process. In the case of consequences, we have "objectives in search of measurable indicators of these objectives," i.e., the attributes of consequences. In the literature on decision theory, this is called an *objectives hierarchy* (Keeney and Raiffa, 1976, pp. 41 - 42). We do this hierarchical inferential process for each goal we have identified. In the process, we can identify many measurable indicators for every goal. In short, we can easily generate many consequence attributes. For example, consider the goal G_9 of taking actions that are politically acceptable to our citizens. How will we generate measurable indicators of the extent to which any of our consequences will reach this goal? What we must do is to generate a collection of attributes for each goal that will apply to all of the consequences we are considering. This does not sound like a very easy task, and it isn't.

The *bottom-up approach* is not any easier, and there are several ways in which this can be done. One way is to have a decision maker compare pairs of possible actions being considered, saying why they prefer one over the other. In the process the decision maker lists the reasons why this preference was asserted. This list of reasons for these preferences is then used to identify consequence attributes.

Another version of the bottom-up approach involves comparing pairs of consequences, say $C_{1,3}$ and $C_{2,2}$. The decision maker is asked to say which one he or she prefers and then lists reasons for this preference. After all pair-wise consequence comparisons have been made, the listing of reasons for these preferences is examined and consequence attributes are identified. This could be a very tedious process since for n consequences there would be n(n-1)/2 pair-wise consequence comparisons to make.

The attributes we will list have been generated by reflecting on the goals we listed in Section 2.1. We listed 10 goals and we provide one measurable indication of each one. There will certainly be more than one indication of the extent to which any of these goals is being met. And there are other possible goals that others may wish to assert. Once again, these attributes apply to each one of our 12 consequences.

A_1: Estimate of how many lives will/could be lost if this consequence occurs.

A_2: Assess the number of threats of hostile actions against each other that are issued by countries affected if this consequence occurs.

A_3: Estimate the extent of discord and hardship among Purple's citizens if this consequence occurred.

A_4: Count the number of developments and tests of nuclear weapons on the part of countries seeking to acquire nuclear weapons or the extent of their nuclear weapons capabilities if this consequence occurs.

A_5: Estimate the number of deaths and injuries among our military and civilian personnel that are likely to result if this consequence occurs.

A_6: Assess our financial costs if this consequence occurs.

A_7: List the reactions of the leadership of other countries, and the international press, regarding attitudes toward us if this consequence occurs.

A_8: Assess the reactions of our military leadership if this consequence occurs.

A_9: Assess the probable reactions of our citizens if this consequence occurred.

A_{10}: Estimate how much the flow of goods from the region of Purple will be reduced if this consequence occurred.

As one can see, all we have done here is to list one *measurable* indication of the values associated with each of the ten goals or objectives we have identified. So, at this point, we have identified ten attributes of the twelve decision consequences we have considered in our Purple example.

The next step is for the decision maker to judge "how much" or "how little" of each these 10 attributes each consequence has.

2.6 Measures and Judgments for Consequence Attributes

2.6.1 Utility or Value?

We must start here with an important terminological distinction. The two terms *value* and *utility* are not synonymous. The difference concerns the numerical scale properties we assume our judgments of the worth of consequences to have (Clemon, 1996, pp. 77 - 82).

Without going into the long history of judgments of worth in fields such as economics and psychology, the term *utility* is used with reference to peoples' worth judgments that are assumed to have *interval (cardinal) scale* properties, equal intervals but no true zero point. For example, we say we like A to degree 0.8 and B to degree 0.6. Then we say that we like C to degree 0.6 and D to degree 0.4. What this is taken to mean is that we prefer A over B by the same amount we prefer C over D. On interval or cardinal scales we can talk about utility differences but not utility ratios. With interval or cardinal scales, we are permitted to make any linear transformation of an original utility judgment that will preserve its properties.

However, suppose our worth judgments have only ordinal properties. Let the symbol ä indicate the judgment "better than." Suppose we say the following: A ä B ä C ä D. In words, we say "We like A better than B, B better than C, and C better than D." We could use a sequence of four numbers here to indicate our preference, such as 1 ä 2 ä 3 ä 4, but the only thing these numbers will indicate is our ranking of these four things as far as our preference is concerned. These numbers will not indicate any differences or ratios. We could use any four other increasing numbers to indicate this ordering; e.g., 12 ä 16 ä 25 ä 112. The reason is that the only thing these numbers indicate is our ranking of these flour things. In much current literature, worth judgments on *ordinal scales* are said to be *value* judgments.

Given the ordinal scale of probabilities used by Disciple-CD (see Table 2), the judgments of the worth of the consequence attributes in our Purple example can only be ordinal in nature and so we should refer to them as *value judgments*.

There are several classical methods for assessing a decision maker's values or utilities (Clemon, 1996, Section 3). For example, in a decision about which house to purchase, we consider such things as:

driving time to work from house (measured in hours/minutes), cost of the house (measured in $$), property taxes (also measured in $$), and size of house (measured in feet^2). We can build a value curve showing the house purchaser's value judgments V(x) for various driving times (x). In this case, the shorter the driving distance to work, the stronger the preference for it. Unfortunately, most of these methods will not be applicable in our Purple example because our attributes have no natural scales. Take A_1 and A_5, for example, involving numbers of deaths. The best we could do is to define fuzzy intervals like "low," "medium," and "high," as shown in Figure 4.

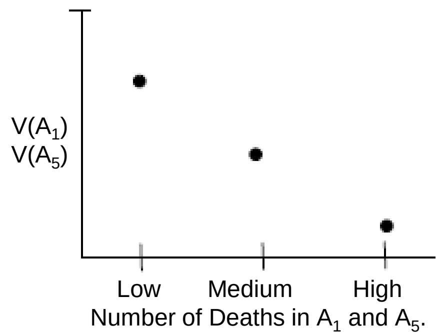

Figure 4. Fuzzy intervals for the attribute "number of death."

Placing a value on human lives/deaths is about the most difficult judgment we could ever think of. Our problem is compounded in this instance because we have only described the three fuzzy ranges of numbers of deaths and would face the questions: How low is "low"? What is a "medium" number? and How high is "high"? In this situation our value judgments can only be weakly ordinal. Of course we would prefer a low number of deaths. The same problem arises in our other attributes where we can only identify fuzzy ranges of measured attributes.

Notice again that using such fuzzy values is consistent with the Baconian-Fuzzy approach used by Disciple-CD to assess the likeliness of the hypotheses.

2.6.2 Consequence Attribute Level Judgments

Below is an example of attribute level judgments for a particular consequence in our Purple example. We have chosen the following consequence to illustrate the process:

$C_{2,3}$: In this case we decide to suspend plans for military actions against Purple and focus on non-confrontational measures. But it happens that Purple is "now" in the act of developing several nuclear weapons having a yield of no more than 20 kilotons each and, within a year's time, will be able to missile launch these weapons at targets within a 1500-mile range. If we discovered these developments, we would have to consider revising this decision.

Here are some examples of observed states or levels of the 10 attributes for consequence $C_{2,3}$. The attributes are those defined above in Section 2.5.

A^*_1: This action would not "now" result in the loss of lives. But some lives might still be lost as a result of accidents during Purple's developing and testing of nuclear weapons. So, we might say that the level of lives lost "now" is "low," if there were no other military actions.

A^*_2: Threats of hostile actions among countries in the Purple's region might "now" increase if it became known that Purple was close to having an arsenal of nuclear weapons, when it also appears that we will not take any military action against Purple if Purple does have an arsenal of nuclear weapons in quite a short time.

A^*_3: The level of discord and distress among the citizens of Purple might "now" decrease as a result of our intended action here, but probably not all that much since Purple will continue to devote its limited resources to the development of nuclear weapons.

A^*_4: If we decide to suspend plans for military actions against Purple and focus on non-confrontational methods, this might "now" encourage other countries to further their developments of nuclear weapons, and other countries to begin developing such weapons. Our action here might be interpreted as a lack of resolve to prevent the spread of nuclear weapons by any means possible.

A^*_5: Our action here would not "now" result in any loss of lives, associated with Purple, of our military and civilian personnel. However, our action here might increase the loss of lives in future if we find we have to respond militarily to Purple's growing nuclear arsenal.

A^*_6: Our action here would "now" decrease the financial cost of dealing with Purple on the theory that diplomacy is cheaper than military action.

A^*_7: This action might "now" win us favorable responses from a number of countries who will applaud our efforts to find a peaceful resolution to problems associated with Purple's nuclear intentions.

A^*_8: This action would not "now" gain the endorsement of our military leaders. They would say that, if we took the non-confrontational action and suspended plans for a military response, it would be difficult to get these plans started again and might be too late if Purple develops its nuclear arsenal in a short time.

A^*_9: Gauging the response of our public to this consequence is probably the most difficult of any of these attributes. Any action taken "now" by our leaders will be disfavored by some persons. We are going to judge that a slim majority would favor the action here because of their unwillingness to see us involved in a war.

A^*_{10}: We are going to say that this action on our part would not "now" result in any decrease in the flow of goods from Purple's region. This action on our part is not at all threatening to the Purple leadership and would give them no excuse to disrupt the flow of goods.

Well, we would all be quite at home arguing about the suitability of the above judgments concerning the levels of the ten attributes for this consequence $C_{2,3}$**.** The same would be true if we had done these judgments for each of the other eleven consequences in our example. But there is something else to notice: all the judgments of these levels have been "fuzzy" in nature.

2.6.3 Combining Value Judgments for Attributes of Consequences

In our Purple example, each of our 12 consequences has 10 measured attributes. The result is that we have 120 value judgments to make, which seems a quite large number unless we can simplify the process. What we wish to end up with is a single value assessment for each of our twelve consequences. It is customary to refer to this single value assessment as being the *composite value* of a consequence. This composite value is obtained by aggregating the value assessments for each of the 10 measured levels of the attributes of that consequence.

The conventional approach to obtain the composite value $V(C^*_{i,j})$ is to assume that the attributes of $C^*_{i,j}$ are all mutually independent and to use a linear combination rule, based on the weights of the attributes:

$$V(C^*_{i,j}) = W_1V(A^*_1) + W_2V(A^*_2) + \ldots + W_mV(A^*_m) + \ldots + W_rV(A^*_r)$$

For different forms of attribute non-independence, a variety of nonlinear rules are possible (e.g., Keeney, 1974). For simplicity, attribute independence is sometimes assumed even when this assumption is not justified.

There are several problems with that approach; the first concerns the independence assumption. The described linear model for determining $V(C^*_{i,j})$ is truly simple but it takes no account of any preference interactions among the attributes A^*_k. It is true that we could change this model to incorporate a few obvious preference interactions. But, when the number of attributes is large, we would have no hope of capturing the many possible interactions that might exist. For example, how much of A_3 the decision maker is willing to give up to get more of A_1 may depend on how much of A_6 and A_{10} the decision maker has. What it comes down to is that our attributes may be dependent in some very complex ways. So, in theory, when we have r attributes, there are $N = 2^r – (r + 1)$ attribute combinations to consider. In short, we have an exponential explosion on our hands and the size of N becomes very important. There are ways of dealing with attribute non-independence for small numbers of attributes, but not for large numbers. The usual strategy is to assume lots of independence among the attributes. An additional point here is that decision theorists have distinguished between different forms of attribute independence (Keeny and Raiffa, 1976; Clemon, 1996, pp. 579 - 580).

But there is another problem so often overlooked in discussions of models like this ones, and it occurs in our Purple example. Many decision theorists have argued that assessments of the value of a consequence are always relative matters. The value of consequence $C_{i,j}$ is always relative to the value of other consequences identified in the same decision. What this says is that consequence values are never absolute, intrinsic, or have immutable worth; they are always relative to the values of other consequences that might occur in the decision at hand. The same can be said about the values we place on individual attributes of any consequence; they are always relative to the values of other attributes of the same consequence. Therefore, if assessments of value are only indications of ordinal preferences, then it makes no sense to add up these $V(A^*_k)$ to obtain $V(C^*_{i,j})$, since these numbers only indicate rankings and have no units.

This relative value idea suggests a way of doing consequence value assessments that may be simpler, depending on the number of attributes in the consequences. This idea is described in the following.

What we need to do is to be able to judge the rank ordering of the 12 multi-attribute consequences $C_{i,j}$ (Section 2.4) in terms of how consistent each one is with our stated ten goals or objectives. Each of the 12 consequences has 10 attributes related to these goals. What we did was to rate the consistency of each attribute for every consequence (120 judgments) on a six-point scale where:

- 3 = High consistency (with the corresponding goal)
- 2 = Medium consistency
- 1 = Low consistency
- -1 = Low inconsistency
- -2 = Medium inconsistency
- -3 = High inconsistency

We then took the sum of the 10 consistency scores for each of the 12 consequences and then assigned each consequence a rank depending on its consistency score. Notice that we could have also used a weighted sum of the consistency scores, by considering the ranking of the goals stated in Section 2.1. The result is shown on Table 4, where the best consequence is $C_{1,1}$ (rank 1), and the worst is $C_{3,4}$ (rank 12).

Table 4. Relative value assessment of options in the Purple example.

	H_1	H_2	H_3	H_4
O_1:	$C_{1,1}$ 1	$C_{1,2}$ 3	$C_{1,3}$ 7	$C_{1,4}$ 8
O_2:	$C_{2,1}$ 2	$C_{2,2}$ 4	$C_{2,3}$ 5	$C_{2,4}$ 6
O_3:	$C_{3,1}$ 9	$C_{3,2}$ 11	$C_{3,3}$ 10	$C_{3,4}$ 12

We can now finally provide a strategy for selecting options in the Purple example. This is discussed in the next section.

3 Selecting an Option

The conventional approach to selecting an option is based on a view that goes back centuries to the time of Daniel Bernoulli and was more recently rescued by the decision theorists Von Neumann and Morgenstern. *Their axioms say that the value of an option or action is equal to the expected value of its consequences* (Von Neumann and Morgenstern, 1944).

Here is precisely where answers to our two basic questions get combined:

- *"What is at stake?"* [indicated by the $V(C^*_{i,j})$ values of the consequences], and
- *"What are the odds?"* [indicated by the P_j probabilities of the hypotheses or states of the world].

What we must determine is $SEV(O_i)$, the *subjective expected value* for all of the n options/actions being considered. But these n options are taken in the face of uncertainty, where we have some number of hypotheses H_1, H_2, …, H_j, …, H_k, …, H_t, representing possible states of the world. We need to put some restrictions on these hypotheses. First, we must assume that they are mutually exclusive and exhaustive; i.e., only one of these hypotheses must be true. Second, we must have subjective posterior probability values $P_j = P(H_j|E^*)$, where E* represents the combined evidence we have on these hypotheses. Third, these P_j must sum to 1.0 in conformance with the Kolmogorov axioms forming the basis of Bayes' rule. Under these conditions, the expression for determining $SEV(O_i)$ is the same one that applies in considering any random variable X in probability theory, $SEV(X) = \Sigma P_j X_j$, where the sum is taken over all possible values of X. So, what this all means, on this interpretation, is that:

$$SEV(O_i) = \sum_{j=1}^{t} P_j V(C^*_{i,j})$$

The process just described allows us to specify a decision rule that many say is optimal: *Choose the option O_i whose $SEV(O_i)$ is largest*. But this process really gets complex when we have masses of evidence and complex consequences, each having many attributes. Don't forget that we have the task of determining the probability of every state of world and the composite value of every consequence.

However, for several reasons discussed below, this conventional approach will not work in our case, especially in our Purple example. First, our four Purple hypotheses (see Section 2.3) are not exhaustive by any means. Because providing an exhaustive set of hypotheses is a very complex task, in the developed computational theory of intelligence analysis and Disciple-CD we do not require the hypotheses to be exhaustive. Second, in evidential reasoning, we do not require users to assess ordinary probabilities. In all of our work so far we have allowed all of our users to provide verbal assessments of uncertainty relying on Fuzzy and Baconian ideas. So, as far as the odds question is concerned, we cannot provide answers that would be compatible with the SEV approach.

Our approach to selecting an option is directly compatible with what we have done so far in Disciple. As shown in Table 4, we have three options and four hypothesized states of the world, so we have twelve consequences, ordered from the best ($C_{1,1}$) to the worst ($C_{3,4}$). So now we are in a position to select an option or course of action. The decision rule here is simple. It says:

Select the option having the highest consequence ranking under the most likely hypothesis.

Here is what our decision rule says with respect to Purple example in Table 4:

- If either hypothesis H_1 or H_2 is most likely, select option O_1.
- If either hypothesis H_3 or H_4 is most likely, select option O_2.

Observe that the options O_1 and O_2 both dominate the option O_3, so we could "now" remove it from consideration. But remember that "now" floats over time and we could reinstate O_3 at any time, pending further information.

4 Justification of the Decision Rule

We will first summarize how we thought of the option selection rule and then offer some comments in its defense. As we have mentioned, Disciple offers a distinct advantage to any user who has both time and resource limitations: It allows the user to "drill-down" to levels consistent with these two constraints. What we have already demonstrated with Disciple-LTA (Tecuci et al., 2008b), TIACRITIS (Tecuci et al., 2010a,b; 2011), and Disciple-CD (Tecuci et al., 2013), are fairly shallow "drill-down" levels in probability assessments based on collections of evidence. For example, we allow the user to employ min/max strategies for assessing the probability of interim and final hypotheses and do not require users to do any full-scale and precise Bayesian or other methods that would require very large numbers of probability assessments. What we have done in this paper is to suggest fairly shallow drill-down levels concerning methods for assessing the composite value of the consequences of options/actions being considered in a decision task. What we needed to do was to come up with a rule for combining these "shallow drilled-down" probability and value-related assessments that would be reasonable and defensible in selecting a course of action. Here is how we came up with this rule.

First, there is a decision rule for selecting a course of action that completely ignores the probability dimension of decisions (Chernoff and Moses, 1959, pp. 8 – 9). That is, this rule does not involve any drilling of probability matters at all, shallow or in depth, but it gives the decision maker the option of choosing whether to be pessimistic or optimistic in his or her views about the consequences of choices. For *pessimists* the rule is called *minimax* and it assumes that, whatever course of action a decision maker chooses, nature will contrive to be in a state, specified by a hypothesis, that is the least favorable to the action the decision maker selects. The designation minimax says that a decision maker should choose the action for which the minimum loss is maximized. In other words, the decision maker should choose the action whose worst consequence is the least risky or unfavorable. We have read accounts saying that this rule is wrongly named, saying that it should be called *maximin* because it says "maximize the minimum loss or risk you face in your choices." For *optimistic decision makers* the rule is *maximax* which says that whatever course of action the decision maker chooses, nature will contrive to be in the state most favorable to this action. So on this optimistic view, the decision maker should choose the action having the most favorable consequence. This is maximizing the maximum gain that could be achieved by the decision maker's choice.

Now, it seems clear that the two rules just mentioned would not make any sense regarding strategies for decisions made in the face of uncertainty. We have said that choices under these conditions require us to answer two very complex questions: *"What is at stake?"* and *"What are the odds?"* The minimax and maximax strategies ignore the second question completely. These strategies do not, however, relieve us of the task of answering the stakes question. But, employing either one of them would lead to eventual ruin of decision makers in business, military, governmental, and a variety of other contexts. There is work suggesting that these strategies only make sense in certain game-theoretic situations. Our work on decisions aids or support systems assumes that the decisions of interest are made in the face of uncertainty and so we must try to answer both stakes and odds questions.

So we asked ourselves the question: *What is the minimal answer we could provide regarding the odds question?* Remember that we need *hypotheses* as well as *options* in order to identify *consequences*. Suppose we have t hypotheses: H_1, H_2, … H_j … H_t. Methods that we have already incorporated in Disciple allow us, at the very least, to rank these t hypotheses in order in terms of their relative likeliness. These methods allow users to assess these rankings under conditions in which they do not have the time or the evidential resources to provide more extensive probabilistic analyses. So, our Disciple methods do allow users to say which of the t hypotheses is most likely, given the evidence we have considered; which hypothesis is second most likely; and so on. A bit later we will consider cases in which two or more hypotheses seem to have the same degree of likeliness.

Now, our next, even more complex, task was to first consider what would be the minimal requirement for judging the relative value of the *patently multi-attribute consequences* of the decisions our intended users will face. Suppose we have n options or actions we are considering in the face of the t hypotheses we are also considering; so we would have n x t possible consequences. We thought, once again, that the minimal requirement would be user's ability to stand these n x t consequences in rank order as far as their judged relative value is concerned. But we now faced the question: *What considerations are absolutely necessary to make these judgments defensible?* The next question becomes: *How can we make these necessary considerations feasible for users having time and resource limitations.* In other words, what is the drill-down level we can employ, with the assistance of Disciple-CD, to make these relative value judgments defensible? These are matters we considered in some detail in the previous sections. Below is a summary account of what we proposed, justifying the value assessment procedures we have provided. The considerations we will mention are those arising in nearly every known account of decision analysis.

(1) The first thing necessary is a statement of objectives or goals to be served in the decision being considered. Our procedures start with the user's specific statements of goals/objectives, recognizing that they will be to varying degrees inconsistent or conflicting, particularly when there are multiple stakeholders involved in the decisions.

(2) The second thing necessary is an account of measurable indications of the extent to which goals are being met in the choice at hand. These measurable indications, deduced from the goals, form the essential *attributes* of the n x t consequences of the choices at hand. Our methods do require identification of the attributes of all n x t consequences being considered, including consideration of the possibility that collections of these attributes can be dependent in various ways. These methods also require a decision maker to state in words each consequence and its major elements. We note here the possibility that any goal might have several attributes associated with it. In the example we gave, for the sake of simplicity, each goal had only one attribute associated with it.

(3) The next matter requires us to state the possible levels or states of each identified attribute. It is appropriate to treat each attribute as a variable or value dimension. All attributes have various states or levels, some of which are to varying degrees consistent with one of our goals and some of which are inconsistent with this goal. The methods we have proposed require the user to first identify the levels or states of each identified attribute.

(4) The decision maker must specify the relative importance of each attribute. This can be done at various stages in the analysis and we have listed it here.

(5) This matter concerns identifying how much of any attribute A does the consequence C actually have. This is mainly where the requirement for measurable attributes comes in. Consequences may of course differ considerably regarding how much of any attribute they have. In other words, different consequences will have different levels or states of any attribute, meaning that they will differ in the degree to which they are consistent with the goal from which an attribute has been deduced. Our methods provide a scale that allows us to assess, for every level of an attribute of a consequence, the degree to which this level is consistent with respect to the goal underlying this attribute.

(6) Now we come to the important task of having the decision maker assess the value associated with having the identified levels of an attribute. This is accomplished by using the scale we have devised that indicates the degree of consistency or inconsistency with a goal for each attribute. This is a six-point scale with three gradations of consistency and three gradations of inconsistency with the goal underlying the attribute.

(7) Next is the necessity for determining the composite value of any consequence when the values of all of its attributes, together with their relative importance, are taken into account. Our methods allow a decision maker to do this under various assumptions about attribute independence. The example we gave assumed that the attributes of consequences were independent. Under this independence assumption, we simply sum together the importance-weighted values of the attributes for each consequence. Now, what results are the composite values for each of the n x t consequences. This allows us to rank-order these consequences in terms of their value. Keep this in mind for the moment.

(8) Now, the final thing necessary is for us to consider is the composite value of each of the n options we are considering, if these composite values can be determined. Remember that each of these n options will have t consequences, each with some different composite value. Now, what is usually recommended in conventional decision analysis, is to determine the *expected composite value* of each of the options we are considering. These expected values are sums of the probability-weighted and importance-weighted composite values across the t consequences of each option.

Number 8 is where we encounter difficulties. Calculating expected values for any option requires an exhaustive set of mutually-exclusive hypotheses, and probability methods that ensure that the distribution of probabilities across these hypotheses sums to 1.0. Bayes' rule, if it can be determined, guarantees that this will happen. Our first problem is that we may not have an exhaustive set of mutually-exclusive hypotheses and, even if we do, we are not supposing that we can do the calculations required under Bayes' rule. Recall that the best we are able to do is to determine the rank ordering of the relative likeliness of our t hypotheses, based on the evidence we have. So, we can never determine the composite value of any of our options by conventional means.

So, how do we decide which option to select? The answer seems quite simple and makes use of what we are assuming about the judgments users of the Disciple agents are capable of providing. First, we look at the hypothesis that seemed most favored by the evidence we have and that we have ranked as "most probable"; suppose this most favored hypothesis is H_j. Then, noting that we have rank-ordered all of our n x t consequences in terms of their composite values, we next look at these rank-orderings for the n options just under H_j, the most likely hypothesis. First, suppose that these rank-orderings of composite consequences under H_j are all different for the n options. Here's the decision rule: Choose the option having the highest-ranked composite value under H_j, the most likely hypothesis; so, suppose this option is O_i. In choosing option O_i, we are taking both probability and value-related considerations into account and are answering both of the two questions as best we can: *What is at stake?* and *What are the odds?* Our essential justification for this decision rule is that it rests, in particular ways, on every major consideration known to us for assessing the composite value of multi-attribute consequences.

But there is a final thing we must mention. Suppose that there are two or more hypotheses having the same probability rankings. Or, even when there is just one hypothesis having the highest ranking, several composite value rankings for the options under this highest ranking hypothesis are tied or are very close together. The decision maker must examine these situations carefully in order to select a course of action, if such selection must be done right away. If this is not the case, the decision maker can defer the decision and buy more information that allows him or her to either revise the probability rankings of the hypotheses being considered or revise the ordering of the composite values of the consequences.

5 Final Remarks

We have attempted to contrive an example involving a currently important situation involving decisions about what our leaders should do regarding Purple's nuclear weapons aspirations and developments. It would be difficult to imagine a more complex example.

Our major goal has been to show how a Disciple cognitive assistant can help users to perform decision tasks having great complexity. At each step in the process, we first showed how conventional decision analysis would offer certain suggestions. But we then showed how infeasible these suggestions would be, given our assumptions about what human judgments are possible, and the limited time available for these judgments. We have suggested instead strategies that combine probability and consequence value judgments which are compatible with what we have so far developed in our work on Disciple.

The following are some of the features of the envisioned decision-support with Disciple, some of them we believe to be unique to Disciple and not found in any other system.

(1) Disciple is designed to take a decision problem from its inception until actions are selected. But it acknowledges that decisions of interest are dynamic and not static in nature. What this means is that, in many national security matters, changes to the ingredients of decisions become necessary. The world continues to change as we are trying to understand it and make well-informed choices and Disciple is designed to handle this.

(2) No system can automatically perform the imaginative reasoning necessary in the generation or discovery of hypotheses and new evidence. However, Disciple assists the user to recognize the mixtures of reasoning required throughout the process of discovery, which is ongoing during decision-making. In such cases we have mixtures of abductive, deductive, and inductive reasoning at work throughout the life cycle of a decision problem. We think that this is a unique feature of our approach.

(3) Disciple has been trained to assist analysts in performing the evidential reasoning tasks that form the basis for hypothesis generation and testing. Our system comes equipped with substantial knowledge of the properties, uses, and discovery of evidence in probabilistic reasoning. Moreover, it continuously learns from expert analysts how to decompose and assess hypotheses, significantly facilitating the analysis of similar hypotheses. These are other unique

features of Disciple and we are not aware of any system that contains similar knowledge about establishing the *relevance*, *believability*, and *inferential force or weight* of evidence or the various *"substance-blind" forms and combinations of evidence* (Schum et al., 2009a,b).

(4) A major feature of Disciple is that it assists the analyst to construct defensible and persuasive arguments from masses of different forms and combinations of evidence in establishing the three credentials of evidence just mentioned. These arguments will certainly be very complex and can be appropriately called *inference networks*. However, the inference networks of concern are the result of trying to make sense out of emerging masses of evidence. Disciple allows the analyst to employ the Wigmorean method of argument construction that has been known about in the field of law for nearly one hundred years (Wigmore, 1913; 1937; Schum, 1987; 2001). Other decision-support systems do not even focus on the construction of arguments from evidence to hypotheses. Unless these arguments are sound and free of disconnects, inferences about the probability of hypotheses will not be persuasive.

(5) The inferential element of choices made in the face of uncertainty result in probabilistic assessments of the relative likeliness of the hypotheses being considered as relevant and important states of the world affecting the outcome of our choices. These probabilities always result from the aggregation of many other probabilities associated with the ingredients identified in the complex evidential reasoning tasks that have been performed. We all know that there are conventional ways of assessing and combining probabilities in evidential reasoning tasks; Bayes' rule comes immediately to mind. But the use of Bayes' rule in the complex argument structures we find necessary comes with a cost that many analysts will not be willing to pay. There will be enormous numbers of probabilities to assess and many of them might be difficult or impossible to evaluate. Add to this the fact that in most inference about matters of national security, the events of concern are singular, unique, or one-of-a-kind. Consequently, there will be no statistics to support probability assessments. Disciple allows the user to make much simpler probability assessments and combinations using ideas that come from the Fuzzy and Baconian systems of probabilistic reasoning. There are many advantages to this simplified approach. This is another unique element of our system.

(6) Another major ingredient of any decision support system concerns assessments of the worth, value, or utility of the *consequences* of options or decisions. A consequence results from choosing option or action O when the world is in state H, where H is some hypothesis being considered. In virtually any decision of interest in national security matters the consequences of concern will have many *attributes* that influence their worth, value, or utility; in short, we must deal with *multi-attribute consequences*. For various reasons we use the term *value* with reference to consequences. The attributes of consequences come from a careful consideration of the goals or objectives asserted by the decision makers. This is really the first question decision makers should ask: *"What do we want to happen as a result of our decision?"* Several or

many goals might be asserted by the decision makers; there may be disagreements among them about what the goals should be. Another area of discord may be that the decision makers will disagree about the relative importance of asserted goals. Finally, these goals will usually be in conflict to some degree. This means that the decision makers cannot have everything they want at the same time. This will signify the necessity for making *tradeoffs* in the choice process.

So, the attributes of a consequence refer to the extent to which this consequence is consistent with the goals stated by the decision makers. A consequence can be consistent with some goals and inconsistent with other goals, and to varying degrees in each case. One necessary thing is that an attribute be a *measurable* indication of how consistent or inconsistent a consequence is with some goal. One way to look at a consequence attribute is that it is a *measurable value dimension*; each attribute tells us how this consequence measures up to one of the goals that have been set. Assessing the value associated with consequence attributes can be a very complex and time-consuming process. We must first state how many possible states there are of any attribute. Then we must say which state applies to a consequence and assess the value of this state. There is a very large volume of literature on multi-attribute value assessment (e.g., Keeney and Raiffa, 1976; Keeney, 1992), but much of this literature involves considerations of this assessment process that most analysts will not wish to consider. We have devised a simplified assessment method that allows the user to grade the extent to which any consequence is consistent or inconsistent with stated goals. Our procedure here is probably unique.

(7) Finally, as far as value assessment is concerned, we have the task of evaluating the composite value of every consequence. We have presented a simple method that lets us determine this composite value for every consequence being considered. But there is another possible value aggregation. If there are N hypotheses, there will be N consequences for any action or decision. There are various algebraic rules for determining the aggregate value of each action. A conventional decision rule here seems simple enough; it says: "Choose the action having the highest expected value." But this rule rests upon our having probabilistic assessments that meet conventional interpretations of probability, something we will not necessarily have. What we needed is a rule for choice that combines our probability and our consequence value assessments. Our rule for such combination is very simple, it says: "Choose the action having the highest valued consequence under the most likely hypothesis." This accounts for both the probabilistic and value elements but it does so in a very simple way. This decision rule is probably quite unique.

6 References

Boicu, M., *Modeling and Learning with Incomplete Knowledge*, PhD Thesis in Information Technology, Learning Agents Laboratory, School of Information Technology and Engineering, George Mason University, 2002.

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

Chernoff, H., Moses, L., *Elementary Decision Theory*, New York, NY, J. Wiley and Sons, 1959.

Clemon, R., *Making Hard Decisions: An Introduction to Decision Analysis, Belmont, CA, Duxbury Press,* 1996.

Cohen, J., *The Probable and the Provable*, Oxford, UK, Clarendon Press, 1977.

Keeney, R., Multiplicative Utility Functions. *Operations Research*, 22, 22-34, 1974.

Keeney, R., *Value-Focused Thinking: a Path to Creative Decision Making*, Cambridge, MA, Harvard University Press, 1992.

Keeney, R., Raiffa, H., *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, New York, NY, J. Wiley and Sons, 1976.

Marcu, D., *Learning of Mixed-Initiative Human-Computer Interaction Models*, Ph.D. Dissertation in Computer Science, Learning Agents Center, Volgenau School of IT&E, George Mason University, January 2009.

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

Pearl J. *Causality: Models, Reasoning, and Inference.* New York: Cambridge University Press; 2009.

Powell G.M. and Schmidt C.F. 1988. A First-order Computational Model of Human Operational Planning, *CECOM-TR-01-8,* US Army CECOM, Fort Monmouth, New Jersey, August.

Schum, D.A., *Evidence and Inference for the Intelligence Analyst* (2 volumes), University Press of America, Lanham, MD, 1987.

Schum, D.A., *The Evidential Foundations of Probabilistic Reasoning*, Northwestern University Press, 1994, reprinting 2001.

Schum, D.A., Tecuci, G, Boicu, M., Analyzing Evidence and its Chain of Custody: A Mixed-Initiative Computational Approach, *International Journal of Intelligence and Counterintelligence,* Vol. 22, pp. 298-319, 2009a. http://lac.gmu.edu/publications/2009/Schum et al - Chain of Custody.pdf

Schum, D.A., Tecuci, G., Boicu, M., Marcu, D., Substance-Blind Classification of Evidence for Intelligence Analysis, in *Proceedings of the Conference "Ontology for the Intelligence Community,"* George Mason University, Fairfax, VA, 2009b. http://lac.gmu.edu/publications/2009/Schum-OIC09.pdf

Tecuci, G., Disciple: *A Theory, Methodology and System for Learning Expert Knowledge*, Thése de Docteur en Science, University of Paris-South, July 1988.

Tecuci, G., *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory, Methodology, Tool and Case Studies,* San Diego: Academic Press, 320 pages, 1998.

Tecuci G., Boicu, M., Marcu, D., Stanescu, B., Boicu, C., and Comello J., Training and Using Disciple Agents: A Case Study in the Military Center of Gravity Analysis Domain, in *AI Magazine*, 24,4, 2002, pp. 51 - 68. AAAI Press, Menlo Park, California, 2002. http://lac.gmu.edu/publications/2002/TecuciG_Disciple_COG_IAAI.pdf

Tecuci, G., Boicu, M., Boicu, C., Marcu, D., Stanescu, B., Barbulescu, M., The Disciple-RKF Learning and Reasoning Agent, *Computational Intelligence, Vol.21, No.4*, pp. 462-479, 2005. http://lac.gmu.edu/publications/2005/TecuciG_Disciple_RKF_CI.pdf

Tecuci G., Boicu M., and Comello J., *Agent-Assisted Center of Gravity Analysis*, CD with Disciple-COG and Lecture Notes used in courses at the US Army War College and Air War College, GMU Press, ISBN 978-0-615-23812-8, 2008a.

Tecuci, G., Boicu, M., Marcu, D., Boicu, C., Barbulescu, M., Disciple-LTA: Learning, Tutoring and Analytic Assistance, *Journal of Intelligence Community Research and Development (JICRD)*, July 2008b. http://lac.gmu.edu/publications/2008/Disciple-LTA08.pdf

Tecuci, G., Schum, D., Boicu, M., Marcu, D., Hamilton, B., 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,* SIST 4, pp. 1-10, Springer-Verlag Berlin Heidelberg, 2010a. http://lac.gmu.edu/publications/2010/Tecuci-Discovery-in-motion-imagery.pdf

Tecuci, G., Boicu, M., Marcu, D., Schum, D., Hamilton, B., TIACRITIS System and Textbook: Learning Intelligence Analysis through Practice, in *Proceedings of the Fifth International Conference on Semantic Technologies for Intelligence, Defense, and Security* – STIDS 2010b. http://lac.gmu.edu/publications/2010/TIACRITIS-STIDS10-paper.pdf

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

Tecuci, G., Schum, D., Marcu, D., Boicu, M., Computational Approach and Cognitive Assistant for Evidence-Based Reasoning in Intelligence Analysis, *International Journal of Intelligent Defence Support Systems*, Vol. 5, 2013.

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

Von Neumann, J., Morgenstern, O., *Theory of Games and Economic Behavior*. Princeton, NJ., Princeton University Press, 1944.

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

Wigmore, J.H., *The Science of Judicial Proof: As Given by Logic, Psychology, and General Experience and Illustrated in Judicial Trials*, 3rd edition, Little, Brown & Co, Boston, MA, 1937.

Zadeh, L., The Role of Fuzzy Logic in the Management of Uncertainty in Expert Systems, *Fuzzy Sets and Systems*, Vol.11, pp. 199-227, 1983.

7 Acknowledgement

This paper is based on research that was partially supported by the Department of Defense and by George Mason University. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies or endorsements, either expressed or implied, of the Department of Defense or the U.S. Government.