Extending Heuer's Analysis of Competing Hypotheses Method to Support Complex Decision Analysis

February 10, 2005

Marco Valtorta (mgv@cse.sc.edu), Michael Huhns, Jiangbo Dang, Hrishikesh Goradia, and Jingshan Huang

Center for Information Technology TR-CIT05-03

Department of Computer Science and Engineering TR-2005-001

University of South Carolina

Abstract

In this paper, we evaluate the Analysis of Competing Hypotheses (ACH) method using a normative Bayesian probabilistic framework. We describe the ACH method and present an example of how to use it to structure an analytic problem. We then show how to represent the same analytic problem using Bayesian networks and compare the result with that using the ACH method. We discuss how Bayesian networks generalize ACH tables and why the added generality might be important to the analyst for hypothesis management. Finally, we propose an approach for acquiring analytic models that interpret situations and for evaluating hypotheses, thereby combining the strengths of ACH and Bayesian networks.

1 Introduction

In general, an intelligence analysis problem comprises three phases: (1) the collection phase, when analysts collect all the evidence pertaining to the problem; (2) the analysis phase, when analysts evaluate the evidence and generate hypotheses; and (3) the reporting phase, when the analysts finalize and submit their results. In particular, the analysis phase involves the management of hypotheses and the application of prior knowledge. We address both of these topics in this paper.

The Analysis of Competing Hypotheses, ACH, is a method to aid judgment on important issues requiring careful weighing of alternative explanations or conclusions. ACH was proposed by Richards Heuer [Heuer99]. Being an effective process that helps avoid common analytic pitfalls, ACH is particularly appropriate for controversial issues when analysts want to leave an audit trail to show what they considered and how they arrived at their judgment or assessments. A software tool based on ACH, called ACH0, has been implemented at PARC [Pirolli04].

A Bayesian network is a graphical model that encodes probabilistic relationships among variables of interest. Because the model has both a causal semantics and a probabilistic semantics, it is an ideal representation for combining data with prior knowledge (which often comes in causal form). ACH can be extended by representing its matrix as a Bayesian network, enabling Bayes reasoning to be used and dependencies among the hypotheses to be revealed and represented explicitly for more in-depth analysis.

In this paper, we evaluate ACH using a normative Bayesian probabilistic framework. Probabilistic frameworks are normative, because any numerical approach to reasoning under uncertainty that satisfies certain obvious requirements "intended to insure consistency with classical deductive logic and correspondence with commonsense reasoning is isomorphic to probability theory" [VanHorn03]. So, probability theory is "a faithful guardian of common sense," and any other approach to plausible reasoning must be considered an approximation, at best, of probability theory [Pearl88, DeFinetti74]. Bayesian networks represent multivariate probability distributions. Their expressiveness and efficiency make them the decision support systems of choice in situations where uncertainty needs to be modeled [Jens01].

The paper is organized as follows. In the second section, we describe the ACH method and present an example of how to structure an analytic problem using ACH. In the third section, we show how to represent the same analytic problem using Bayesian networks. The fourth section generalizes the example by comparing ACH with Bayesian networks. The fifth section contains a discussion of how Bayesian networks generalize ACH tables and why the added generality may be important. The sixth section describes a new approach to acquisition of analytic models for interpretation of situations and evaluation of hypotheses. The approach combines the strengths of ACH and of Bayesian networks.

2 Related Work

Different analytical methodologies enable analysts to organize and focus their energies during analysis, as follows [Nunn04]: (1) A delphi technique finds consensus from a group of subject matter experts; (2) Formulaic mode is a statistical approach that assigns each course of action (COA) a numeric percentage-based probability of adoption; (3) Probability diagrams are graphic depictions of relationships and activities; (4) Inductive reasoning makes broad assumptions based on known facts; (5) Deductive reasoning takes a known event and breaks it down to determine the exact events. Based on the above methodologies, various analytic tools have been developed to assist analysts in accomplishing their analysis tasks.

By combining the concepts from structured argumentation and ACH, Cluxton et al. [Cluxton04] constructed an information visualization tool for understanding complex arguments. This tool enables analysts to manipulate the hypotheses and their associated inference networks by linking evidence with hypotheses and setting evidence parameters such as relevance and credibility. This tool does not, however, capture the analysts' prior knowledge or provide services such as sensitivity analysis and surprise detection by probabilistic reasoning.

In addition to ACH, which is the most widely used, other techniques have been suggested for analyzing intelligence problems. The following three alternative analysis techniques focus on understanding how to analyze industrial trends and how to add insight to information. Although specialized to industrial problems, they might be adapted to analysis tasks as competitors to ACH.

1. Porter's Five-Forces Analysis

   Michael E. Porter’s Five Forces model [Porter98, QMBA04, Manager01] is an approach to analyzing industrial structure based on five competitive forces acting in an industry or a sub-industry: threat of entry, threat of substitution, bargaining power of buyers, bargaining power of suppliers, and rivalry among current competitors. Based on the information derived from application of Five Forces Analysis, management can decide how to influence or exploit particular characteristics for their industry. However, the model has some limitations in today’s market environments, because it does not take into account new business models and the dynamics of markets [Manager01]. A similar analysis technique, substituting analytical task dynamics for the competitive forces, could make this methodology usable by intelligence analysts.

2. Win-Loss Analysis

   Win-Loss Analysis is a business-to-business research tool that attempts to provide high-quality information quickly and cost-effectively, targeting the specific people that make purchasing decisions [graffgroup]. It impacts the sales process at every point and provides actionable insight from a historical as well as a predictive viewpoint. By obtaining reliable and unbiased feedback from recent sales contacts, sales representatives can refine their techniques, learn how to effectively target a client’s needs and the appropriate decision makers, and place the company in the best possible light [primary intelligence].

   Unfortunately, actionable and accurate win-loss data is difficult to obtain by a single entity. Unbiased, third party win-loss data collection and analysis, a comprehensive methodology for uncovering the root cause of wins-losses, and actionable recommendations to improve the resultant win percentage remain active areas for further development [current analysis].

3. Scenario Planning

   Scenario Planning is a model for learning about the future in which a corporate strategy is formed by drawing a small number of scenarios, i.e., stories about how the future may unfold, and showing how these might affect an issue that confronts the corporation. It works by understanding the nature and impact of the most uncertain and important driving forces affecting the future.

   Scenario Planning is most widely used as a strategic management tool, but it is also used for enabling group discussion about a common future [value based management]. Being a group process that encourages knowledge exchange and development of mutual deeper understanding of central issues important to the future of the business, Scenario Planning’s goal is to craft a number of diverging stories by extrapolating uncertain and heavily influencing driving forces. However, a fairly complex set of attributes might have to be determined in advance, which limits the extensive application of Scenario Planning [MIT77].

3 The ACH Approach and Its Use

One way that some analysts go about their business is via a satisficing strategy, whose principal weakness is the failure to recognize that most of the evidence for the single hypothesis chosen might also be consistent with other alternatives not been refuted. However, simultaneous evaluation of competing hypotheses is difficult to carry out for most people. Fortunately, with the help of ACH, that task is much easier to be accomplished [Heuer99]. The following description outlines the steps taken in ACH.

1. Identify the possible hypotheses to be considered. Make a list of significant evidence and arguments for and against each hypothesis.
2. Build a matrix with hypotheses across the top and the evidence down the side, and analyze the diagnostic value of each piece of evidence with respect to each hypothesis. Refine the matrix and repeat this step when necessary.
3. Draw tentative conclusions about the relative likelihood of each hypothesis by trying to disprove the hypotheses instead of proving them.
4. Analyze the sensitivity of each conclusion in step 3 to a few critical items of evidence, then report final conclusions by discussing the relative likelihood of all hypotheses rather than the most likely one, and identify milestones for future observation that may indicate events are taking a different course than expected.

While there is no guarantee that ACH will produce a correct assessment, it does provide an appropriate process of analysis through which the odds of getting the right answer increase greatly.

Throughout the rest of the paper, we use a fictitious example. We imagine that an analyst who is a specialist on terrorist activities related to the oil infrastructure of Iraq and Iran has to evaluate hypotheses in the Abadan region of Iran. The interest in evaluating the hypotheses is high, because of the recent interception of a message between terrorists. We emphasize that this is a fictitious example, devised to illustrate our techniques.

Question: Will terrorists try to create conflict in Iran by attacking the oil infrastructures in Abadan region?

Hypotheses:
H1: Terrorists will bomb the oil refineries in Abadan.
H2: Terrorists will bomb the oil pipelines in Abadan.
H3: Terrorists will bomb the oil wells in Abadan.
H4: Terrorists will bomb the oil facilities in Shiraz.
H5: Terrorists will not launch an attack.

Evidence (fictitious for this example):
E1: A phone wiretap on a suspected terrorist cell in Beirut records a discussion about crippling the Iranian economy by destroying oil production facilities within the Abadan region.
E2: The oil refinery in Abadan can produce 0.37 million barrel per day. Oil is transported through pipeline.
E3: the oil refinery in Shiraz can produce 0.04 million barrel per day.
E4: There is an oil pipeline with from Abadan to Basra, which crosses the border. The capacity of this pipeline is over 0.2 million barrel per day.
E5: Historical analysis allows us to conclude that the affected oil industry will cripple the Iranian economy, which will lead to the conflict with its neighbors.

E6: The area near a border is easier for terrorist to infiltrate.
E7: Terrorists prefer a target that is near a road.

The preceding question, hypotheses, and items of evidence lead to the ACH matrix presented in the following table.

Table 1: An ACH Matrix

	H1	H2	H3	H4	H5
E1	+	+	+	-	-
E2	+	+	+	-	-
E3	-	-	-	+	-
E4	+	+	-	-	-
E5	+	+	+	+	-
E6	-	+	-	-	-
E7	-	-	-	-	-

4 Bayesian Network Representation of ACH Tables

Bayesian networks are a space-efficient representation of multivariate probability distributions that exploits independence information and supports the time-efficient computation of posterior probabilities. The expressiveness and efficiency of Bayesian networks make them the decision support systems of choice in situations where uncertainty needs to be modeled [Jens01].

More precisely, a *Bayesian network* [Pear88, Neap90, Jens01] consists of a directed acyclic graph (DAG), called a *Bayesian network structure*, prior marginal probability tables for the nodes in the DAG that have no parents, and conditional probability tables for the nodes in the DAG given their parents. The network and the probability tables define a joint probability distribution on all variables corresponding to the nodes, with the defining property that the conditional probability of any variable *v* given any set of variables that includes only the parents of *v* and any subset of nodes that are not descendant of *v* is equal to the conditional probability of v given only its parents.

From this property, it follows that the joint probability of the variables in a Bayesian network decomposes in a multiplicative fashion; more precisely, if *V* is the set of the nodes in the DAG, the following equality (the *chain rule for Bayesian networks*) holds: $P(V) = \prod_{v \in V} P(v \mid parents(v))$. In turn, this decomposition allows for the very efficient computation of marginal posterior probabilities upon observation of evidence.

We now show that the ACH table of Table 1 can be represented as a bipartite graph, where the nodes are divided into two exhaustive and mutually exclusive sets, corresponding to hypotheses (columns in the ACH matrix) and items of evidence (the rows in the ACH matrix, also called findings). Figure 1 below shows the resulting Bayesian network structure.

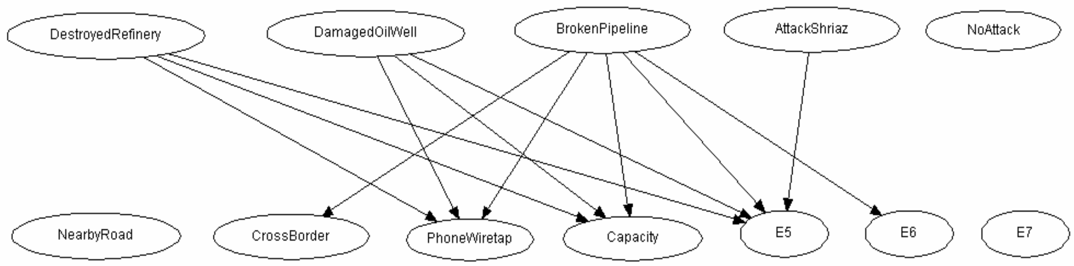

Figure 1: Bayesian network corresponding to the ACH matrix in Table 1

Heuer suggests using a simple linear, additive scoring mechanism to assess the probability of a hypothesis. Such a scheme can easily be incorporated within the Bayesian network framework, as described in [Oles93]. However, as Heuer himself notes, it is sometimes preferable to use probabilities rather than a plus and minus notation. In particular, we observe that it is also possible and preferable to represent the sensitivity and specificity (or "diagnosticity," to use Heuer's term) of items of evidence for hypotheses directly in conditional probability tables. For example, we can represent a situation for which E4 ("CrossBorder") is a moderately sensitive but very specific item of evidence for the hypothesis H2 ("BrokenPipeline") as indicated in Table 2.

Table 2: CrossBorder is highly diagnostic for BrokenPipeline, but only moderately sensitive to BrokenPipeline

P(CrossBorder|BrokenPipeline)	BrokenPipeline = yes	BrokenPipeline = no
CrossBorder=yes	0.7	0.01
CrossBorder=no	0.3	0.99

In his book, Heuer makes it clear that it is very important to specify prior beliefs in order to obtain correct posterior beliefs. The translation of ACH matrices to Bayesian networks ensures that prior probabilities of hypotheses are assessed.

5 Comparing ACH Matrices and Bayesian Networks

By ACH, an analyst explicitly generates all reasonable hypotheses and compares them by analyzing the diagnosticity of a list of related evidence and arguments, and then draws the tentative conclusions about the likelihood of each hypothesis. The BN reasoning service that we have developed within the Novel Intelligence from Massive Data (NIMD) program can be used to enhance ACH by automating some of the analysis steps and by supporting extensions to the representation of uncertainty used by ACH. Here we go through ACH by comparing it with our BN reasoning service. The discussion that immediately follows is organized around the eight steps of the ACH technique as written by Heuer [1999].

1. Identify the possible hypotheses to be considered.

In BN Reasoning, the possible hypotheses are represented by nodes in Bayesian networks. The hypotheses that appear to be disproved and the hypotheses that are simply unproven are both nodes with the near-to-zero probability values in BNs. The difference between them is that the disproved hypothesis nodes have evidence nodes that link to them, while unproven hypothesis nodes don't. We keep the unproven hypotheses until they are disproved by items of evidence (typically contained in incoming messages).

2. Make a list of significant evidence and arguments for and against each hypothesis.

Explicit evidence is represented by the messages and the assumptions or arguments that the analysts made are captured as the prior knowledge and stored in our BN fragments repository.

3. Prepare a matrix with hypotheses across the top and evidence down the side. Analyze the "diagnosticity" of the evidence and arguments--that is, identify which items are most helpful in judging the relative likelihood of the hypotheses.

4. Refine the matrix. Reconsider the hypotheses and delete evidence and arguments that have no diagnostic value.

5. Draw tentative conclusions about the relative likelihood of each hypothesis. Proceed by trying to disprove the hypotheses rather than prove them.

Diagnosticity of evidence is captured by the conditional probability tables residing in the BN fragments. To judge the relative likelihood of hypotheses, we can compute the probability values of each hypothesis after matching the messages with the BN fragments and composing the BN fragments into a situation-specific scenario. The probability representation of the likelihood is finer than the notation of minus/plus or the numerical scale in ACH.

6. Analyze how sensitive your conclusion is to a few critical items of evidence. Consider the consequences for your analysis if that evidence were wrong, misleading, or subject to a different interpretation.

Our program provides a sensitivity analysis service that can be used to evaluate how sensitive the target (hypothesis) is to the provided evidence.

7. Report conclusions. Discuss the relative likelihood of all the hypotheses, not just the most likely one.

8. Identify milestones for future observation that may indicate events are taking a different course than expected.

The tacit knowledge is captured in the form of additional BN fragments. Our program provides a value of information service that recommends which additional pieces of evidence should be collected in order to increase the confidence in the conclusions reached.

6 How Bayesian Networks Enhance ACH Matrices

Since ACH tables can readily be converted to bipartite Bayesian networks, it is natural to ask whether we can exploit this translation. Fortunately, the answer is positive. First, we observe that bipartite Bayesian networks are a special case of Bayesian networks. There are limitations to the expressiveness of bipartite Bayesian networks. First, it is impossible to represent dependency among hypotheses that is not mediated by items of evidence. In other words, in the absence of evidence, one's belief in a hypothesis cannot affect the belief in another hypothesis. This is clearly inappropriate in situations in which a model exists of how hypotheses affect each other. Second, it is impossible to represent dependencies among items of evidence that are present even when the hypotheses are known. Such dependencies would be modeled by introducing intermediate variables between hypotheses and items of evidence. We note that this is a particularly serious issue when trying to model rumors and deception. Third, it is impossible to model context for hypotheses. Therefore, we develop a more complex model for our motivating example, as shown in Figure 2, to overcome all these limitations:

1.) We model a conflict situation, in which context is represented by the two related variables Conflict and AffectedOilProduction.

2.) We introduce the intermediate nodes such as TerroristAction and ThreatLevel to represent the dependencies among items of evidence in Bayesian Network.

3.) We represent the argument, the assumption the analyst made in ACH, as the structure of the BN fragment, instead of nodes in bipartite graph model.

4.) In our model, the hypotheses are related through the context, even in the absence of evidence.

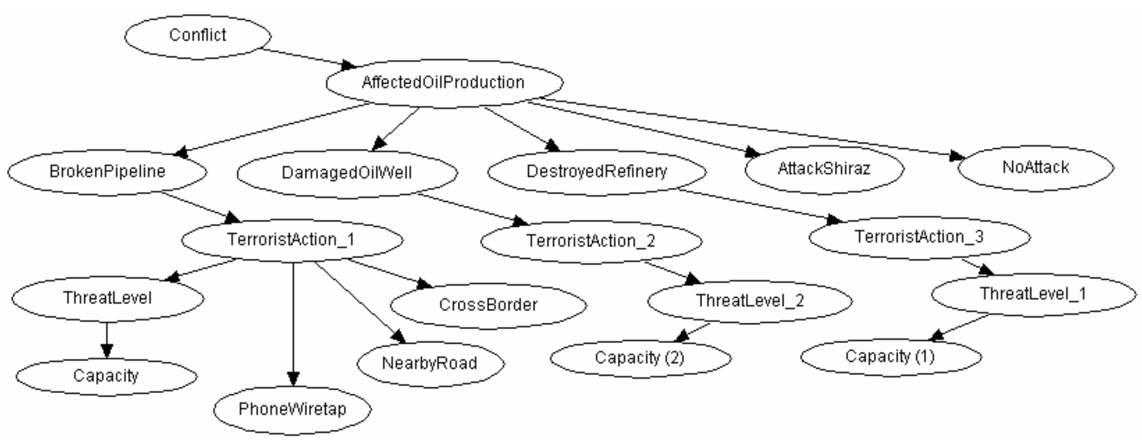

Figure 2: A more complex model for the oil facility example of Table 1 and Figure 1

7 How to Integrate ACH and Bayesian Networks

From an ACH table, we can obtain a bipartite graph by adding a link from the hypothesis node to the evidence node if the corresponding value in the ACH matrix is either + or -. To assign the initial conditional probability values, we define the number of items evidence as *n,* which is 7 in our example, and let the initial probability $P(E_i|H_j)$ be $1/n$ if there is a link from H_j to E_i. The probability of $P(-E_i|-H_j)$ is assumed to be 0.9. We emphasize that this is only an initial assignment.

By creating BN fragments (bipartite graphs and associated conditional probabilities) from ACH tables, we capture prior knowledge, which we then refine by adding context variables and intermediate nodes. Many prior or conditional probability values can be obtained from the statistical information available in databases.

8 Conclusions and Future Work

In this paper we have extended the concept of ACH by incorporating Bayesian reasoning. The ACH matrix is represented as Bayesian network fragments and our reasoning services facilitate a detailed study of the intelligence analysis problem by the analyst. However, we have not yet addressed all of the implications or tested the extension in an operational environment. To perform such testing, we will first develop a tool that generate Bayesian Network fragment from the ACH table. Moreover, our tool will support the analysts in discovering inadequacies in the bipartite Bayesian network obtained from the ACH table. Such support will be provided by a tool that uses sensitivity analysis and a collection of past cases.

References

[Heuer99] Richards J. Heuer, Jr. *Psychology of Intelligence Analysis*. Center for the Study of Intelligence, 1999. (At http://www.cia.gov/csi/books/19104/index.html)

[DeFinetti74] Bruno De Finetti. *Theory of Probability, Volume 1*. Wiley, 1974.

[Jensen01] Finn V. Jensen. *Bayesian Networks and Decision Graphs.* Springer, 2001.

[Neap00] Richard E. Neapolitan. *Probabilistic Reasoning in Expert Systems: Theory and Algorithms.* Wiley, 1990.

[Oles93] KristianG. Olesen and Steen Andreassen. "Specification of Models in Large Expert Systems Based on Causal Probabilistic Networks." *Artificial Intelligence in Medicine*, 5, 3 (June 1993), pp.269-281.

[Pearl88] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufman, 1988.

[Pirolli04] Peter Pirolli and Lance Good. "Evaluation of a Computer Support Tool for Analysis of Competing Hypotheses." Palo Alto, CA: Palo Alto Research Center (UIR Technical Report).

[Porter98] Michael E. Porter. *Competitive Strategy: Techniques for Analyzing Industries and Competitors.* Free Press, 1998

[QMBA04] http://www.quickmba.com/strategy/porter.shtml

[VanHorn03] Kevin S. Van Horn. "Constructing a Logic of Plausible Inference: A Guide to Cox's Theorem." International Journal of Approximate Reasoning, 34, 1 (September 2003), pp.3-24.

[Manager01] http://www.themanager.org/pdf/p5f.pdf

[Nunn04] Matthew J. Nunn "ASAS Master Analysts' support to information operations—analysis", *Military Intelligence Professional Bulletin,* 2004.

(At http://www.findarticles.com/p/articles/mi_m0IBS/is_1_30/ai_n6123727)

[Cluxton04] Diane Cluxton, Stephen G. Eick, and Jie Yun "Hypothesis Visualization", *Proceedings of the IEEE Symposium on Information Visualization (INFOVIS'04)*, pp. 215.4

[graffgroup] http://www.graffgroup.com/winlossarticle.html

[primary intelligence] http://www.primary-intel.com/solutions/winloss.aspx

[current analysis] http://www.currentanalysis.com/intelligence/win_loss.cfm

[value based management] http://www.valuebasedmanagement.net/methods_scenario_planning.html

[MIT77] Bradley, Hax, and Magnanti. *Applied Mathematical Programming.* Addison-Wesley, 1977

Acknowledgments

This work was supported in part by the NIMD Program under the Advanced Research and Development Activity (ARDA). Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the U.S. government. We thank Bob Simpson, Betty Whitaker, Laura Burkhart, Collin Lobb, and Reid McTavish, our collaborators at Georgia Tech in the NIMD project "Case-Based and Bayesian Reasoning for Knowledge Discovery and Threat Situation Analysis." We also thank our Contracting Officer's Technical Representative.