LA-UR- 09-04277

Approved for public release; distribution is unlimited.

Title:	Utilization of Extended Bayesian Networks in Decision Making under Uncertainty
Author(s):	Ed Van Eeckhout, Deborah Leishman, Bill Gibson
Intended for:	2009 MSS National Symposium on Sensor & Data Fusion (NSSDF): Proceedings

Los Alamos
NATIONAL LABORATORY
EST. 1943

Form 836 (7/06)

Utilization of Extended Bayesian Networks in Decision Making under Uncertainty

August 2009

Ed Van Eeckhout, Deborah Leishman, and Bill Gibson

Los Alamos National Laboratory
Los Alamos, NM

ABSTRACT

A Bayesian network tool (called IKE for Integrated Knowledge Engine) has been developed to assess the probability of undesirable events. The tool allows indications and observables from sensors and/or intelligence to feed directly into hypotheses of interest, thus allowing one to quantify the probability and uncertainty of these events resulting from very disparate evidence. For example, the probability that a facility is processing nuclear fuel or assembling a weapon can be assessed by examining the processes required, establishing the observables that should be present, then assembling information from intelligence, sensors and other information sources related to the observables. IKE also has the capability to determine tasking plans, that is, prioritize which observable should be collected next to most quickly ascertain the "true" state and drive the probability toward "zero" or "one." This optimization capability is called "evidence marshaling."

One example to be discussed is a denied facility monitoring situation; there is concern that certain process(es) are being executed at the site (due to some intelligence or other data). We will show how additional pieces of evidence will then ascertain with some degree of certainty the likelihood of this process(es) as each piece of evidence is obtained. This example shows how both intelligence and sensor data can be incorporated into the analysis.

A second example involves real-time perimeter security. For this demonstration we used seismic, acoustic, and optical sensors linked back to IKE. We show how these sensors identified and assessed the likelihood of "intruder" versus friendly vehicles.

1.0 Introduction

For some time now, we have been applying Bayesian Belief Networks (BBNs) to problems involving multisource data fusion. In simple terms, Bayesian Belief Networks process "evidence" to compute probabilities of "hypotheses." For example, some of our problems have involved monitoring of an adversary's actions to determine intent (hostile or benign), or monitoring of a remote facility to determine what types of covert processing might be done there. In these cases the evidence might be extracted from textual intelligence messages acquired from overhead reconnaissance assets or other types of intelligence. In these applications, evidence is costly and risky to obtain and one would want to optimally task the intelligence gathering assets to collect the best evidence to reach conclusions quickly and with reasonable costs. In these problems, the hypothesis nodes in the BBN will likely represent competing alternatives as to what the adversary is really doing. Another type of problem involves near-real-time surveillance for the purpose of threat detection and identification. In such cases, the evidence is extracted from real-time sensor data feeds as well as other types of sources.

The consequences of an incorrect decision can be very high, depending upon the situation at hand. The analyst and decision maker not only want the answer, but want to know with as much certainty as possible. To help with this, we have developed enhancements to traditional Bayesian Analysis which quantify the uncertainty of subject matter expert statements that are themselves probabilistic in nature. We then combine those with evidence from ground, air, or satellite sensors that have established uncertainty bounds. The enhancements we will describe have to do with treating model parameter uncertainty, incorporating evidence uncertainty, determining the optimal evidence to collect next (evidence marshalling), and determining the best asset with which to collect the evidence (asset allocation).

Procedurally, our approach was to select a well-known traditional Bayesian Analysis tool called Netica™ (typically used by researchers), as our Bayesian inference engine. Our enhanced Analysis tool, called the Integrated Knowledge Engine (IKE) wraps around Netica to "extend" traditional Bayesian analysis to better handle monitoring and surveillance problems and to make it more suited for use by intelligence analysts and decision makers.

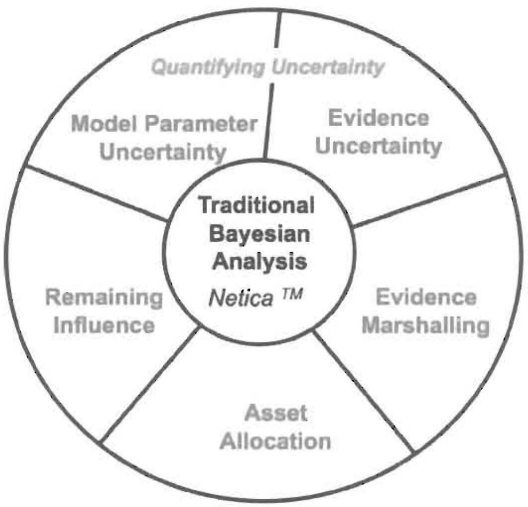

Figure 1. Our "extension" of traditional Bayesian analysis.

2.0 Bayesian Belief Networks

Bayesian Belief Networks (Pearl, 2000; Jensen and Nielsen, 2007) provide a way to conceptualize and model problems which involve trying to reach conclusions based on evidence. Often they are used to try to understand the causal relationships between a set of variables. BBNs have been successfully applied to various problem domains such as medical (diagnosis), judicial (guilt/innocence), and forensics (what happened), to name a few. A BBN consists of nodes and directed links (arrows) connecting the nodes. One can think of the arrow as representing a parent-child relationship. Often in Bayesian modeling, the arrow represents a causal relationship. Hence the rule of thumb – "Parents cause the children."

A node in a BBN represents a variable that can be in only one of finitely many states. For example if temperature is a variable, one could say that the temperature could be hot or cold. Thus the temperature node would be modeled as having two states: hot and cold. A more complicated problem might require that temperature have four states; freezing, cold, warm, hot. We may not know what state the temperature node is in, in which case we say that the temperature node is unknown. If we get some data that tells us the state is hot, we can enter that finding into the BBN and "set" the temperature node's state to hot. This is called entering evidence (or entering a finding).

Often the top-level nodes in a BBN represent the competing alternatives that we are trying to sort out: is the factory making fertilizer, anthrax, sarin, or something else? These nodes are called hypotheses nodes. Often the bottom level nodes in the network represent things we can observe as evidence: is the factory using low, medium, or high amounts of electricity? These nodes are called evidence nodes. Given the evidence that we have entered into the BBN, we want the BBN to calculate the probabilities of the states of the hypothesis nodes to reach a conclusion such as: The probability that the factory is making fertilizer might be 88%, anthrax 7%, sarin 3%, and other 2%. Several algorithms (Huang and Darwiche, 1996) have been developed that can perform this Bayesian Inferencing in a practical and useful manner – provided the network is a BBN.

In order for a network to be a BBN it must satisfy two conditions: (1) It must be an Acyclic Directed Graph – acyclic in the sense that there are no loops in the graph – i.e. a parent may not be the child of one of it's descendents, and (2) The network must satisfy the Markov condition – if the state of all the parents of a node are known, then the state of that node is influenced only by it's descendents. These conditions simplify the problem enough that it can be solved by Bayesian Inferencing algorithms. Condition (1) keeps the algorithms from encountering an infinite loop, and condition (2) allows the model builder and the algorithms to worry only about the immediate relationships between a node and it's parents.

Figure 2 shows such a network: this network was constructed to determine which process is occurring (e.g., fertilizer, anthrax, sarin, and "other") if a facility is operational. To construct a BBN one must provide a table of conditional probabilities (CPs) for each node. These CPs represent our answer to the question: Given the state of all my parents, the probabilities for my states are the following. For example, the CP Table (CPT) for the e43 node (Grd Sensor), is shown. The CP shown for the highlighted cell means: Given that e35 (sub7) is "true," the probability of e43 (Grd sensor) is "true" is 75%. Because we are reasoning from cause to effect (from parent to child), a subject matter expert who is familiar with the system being modeled can easily define these conditional probabilities and populate the table. Alternatively, the CPs can be derived from data, if available. It is much harder to reason from effect to cause, but this is precisely what the Bayesian inferencing algorithms do for us, as shown in Figure 3. We enter the observed effects into the BBN as evidence, and the BBN performs inferencing to infer the probable cause.

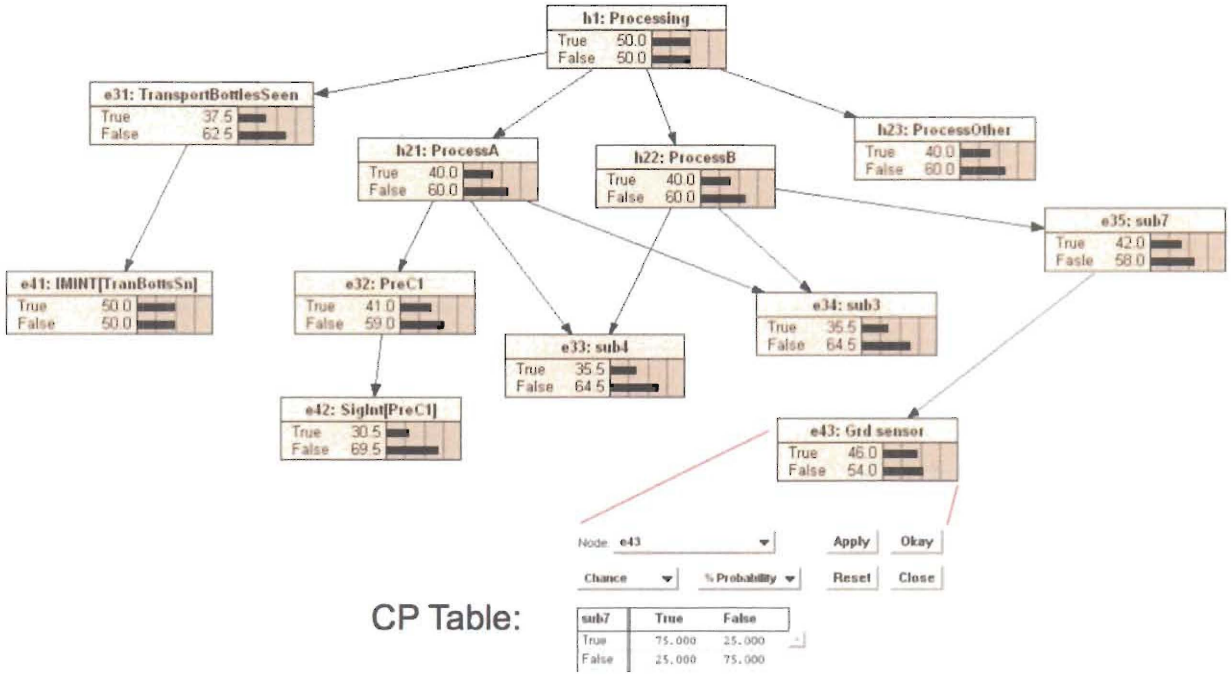

Figure 2. An Example Bayesian Network for Facility Processing.

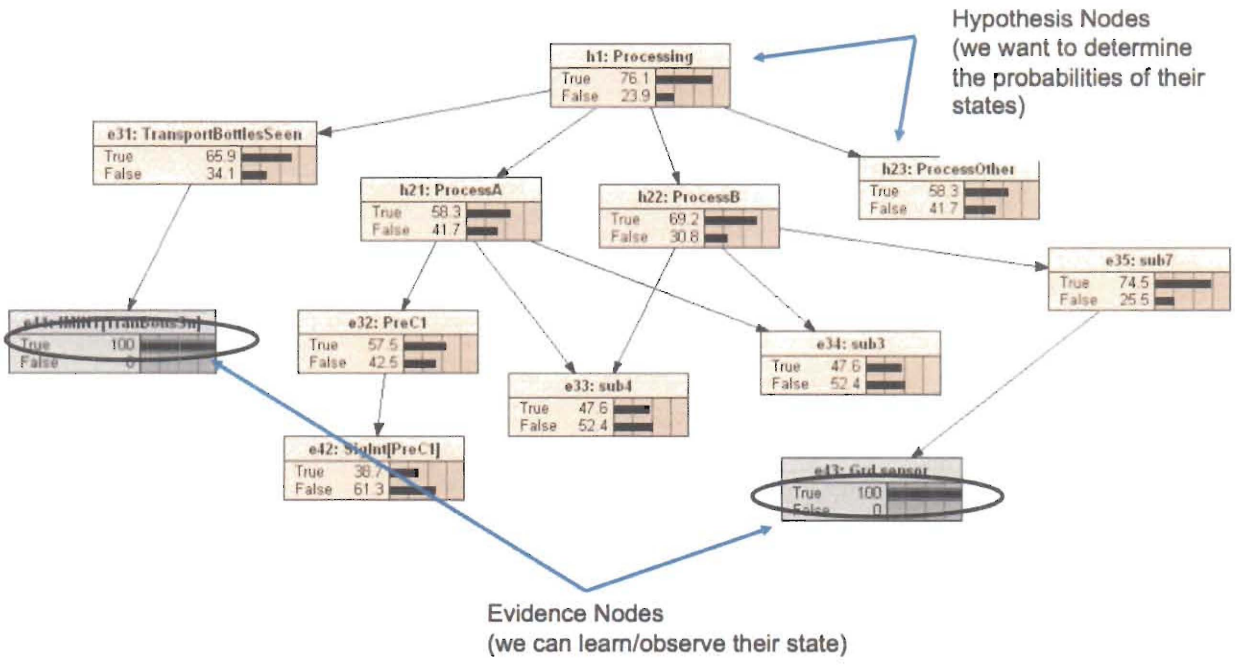

Figure 3. Example Evidence Observed in Processing Network.

3.0 The Integrated Knowledge Engine (IKE)

The Integrated Knowledge Engine enhances the traditional Bayesian Analysis provided by the Netica engine, by providing the following additional major capabilities:

- Analysis (Inferencing with Uncertainty),
- Evidence Marshalling (with Uncertainty),
- Asset Allocation, and
- Remaining Influence.

IKE's flexible graphical user interface can be easily changed, so each new application of IKE may have it's own look and feel, while the underlying classes that implement the core capabilities remain unchanged. Further details on these capabilities are described in Gibson and others (2009).

Figure 4 shows the graphical interface for our simple facility processing example discussed above. The hypotheses and their current state based on the evidence are shown in the middle. The Evidence is shown on the right, and the inset shows the various forms in which evidence can be entered. In this case, we are 80% "confident" that transport bottles have been observed. This results in uncertainties being propagated in terms of second order uncertainties on the hypotheses probabilities (shown as + and – values after the mean in the hypotheses window) as described by Izraelevitz and others (2007). These uncertainties are obtained by using a Monte Carlo simulation by independently setting the Netica "true/false" or "hi/medium/low" bins in the appropriate percentages over at least 1000 simulations.

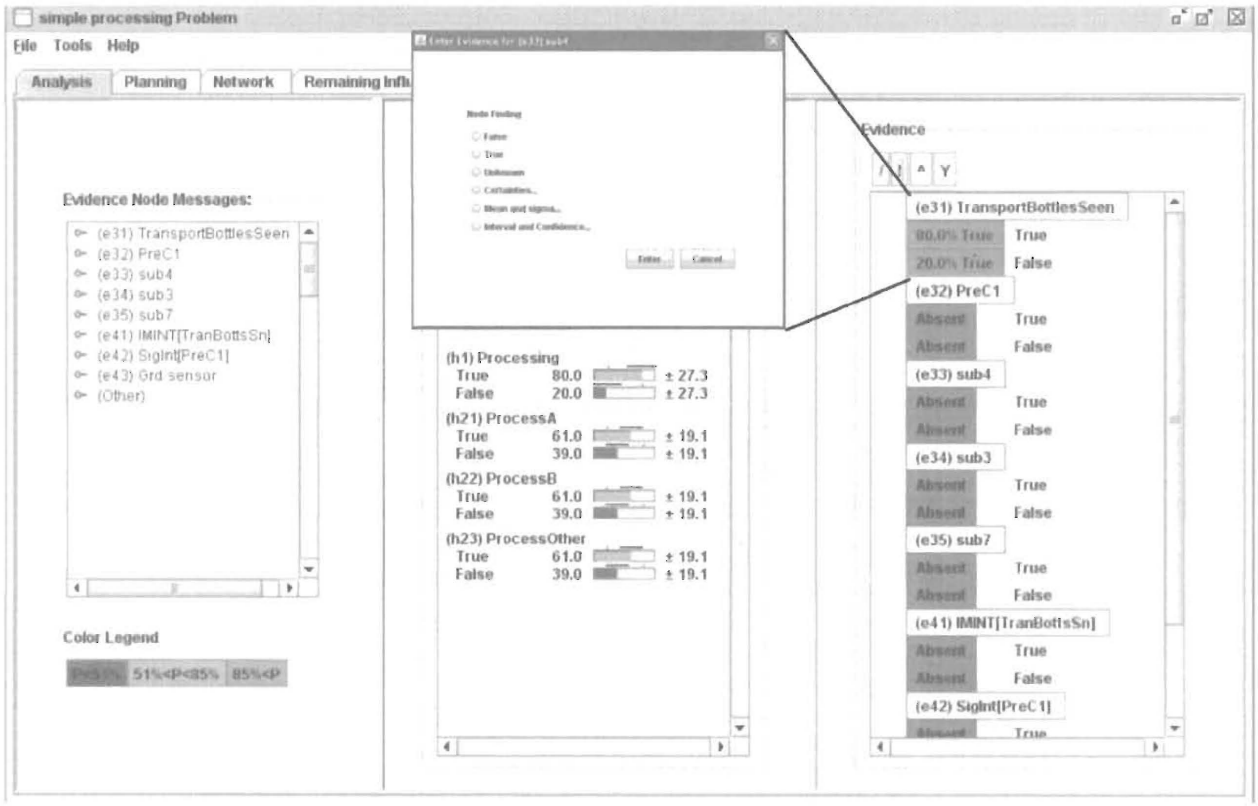

Figure 4. IKE Interface.

In the simplest configuration, IKE is used like a hand calculator with the analyst manually entering evidence as it becomes available from intelligence reports or sensor data. In some applications, intelligence or sensor reports arrive at IKE only to be placed in the mailbox for an evidence node (some

process has routed the report to the appropriate node). The analyst reviews the report manually and decides whether to enter evidence. In other near-real-time applications, special IKE evidence messages are automatically generated from the multisource data streams and set their evidence into IKE automatically. We have found that viewing incoming data streams as sources of discrete evidence in a Bayesian Belief Network provides a framework (and a simple architecture) for integrating diverse types of input data for knowledge discovery.

4.0 Processing Facility

To expand upon the simple situation defined above, a more detailed model was created to monitor a facility whose location was known, but if it started processing, we didn't know what it might be processing. Various processes could be called hypotheses A, B, and C in Figure 5 below.

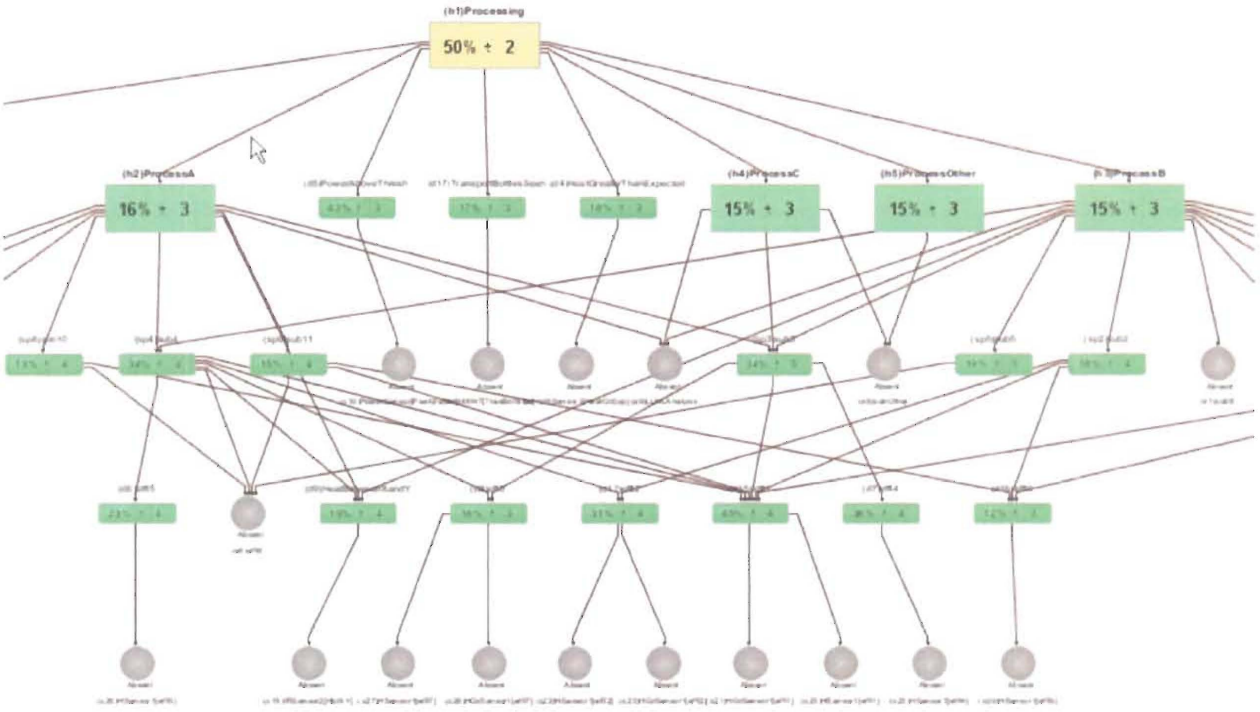

Figure 5. Part of our Network for a Processing Facility.

Each process is linked to subprocesses with certain potential signatures (in terms of heat, effluents, power, etc). So at the very bottom are evidence/signatures that indicate subprocesses are interest. If these signatures are present, then we have a certain confidence that one of the subprocesses (and thus eventually hypotheses) is true or false. Note this particular network is binary: all nodes are true or false, so only the probability of "true" is displayed.

To illustrate the process, Figure 6 shows a possible timeline for receiving evidence. This can then be displayed on the IKE screen as evidence is added with time (Figure 7). The hypotheses are shown in the center of the screen; the evidence is shown coming in directly from sensors on the right or messages on the left; and the current probability and uncertainty is shown with the hypotheses. The calculations of these values are a result of the Bayes net conditional probabilities for each node, the prior probabilities, and the evidence. In this particular example, the entire network has been "triggered" by a power anomaly. Time after the power anomaly is shown in the lower left.

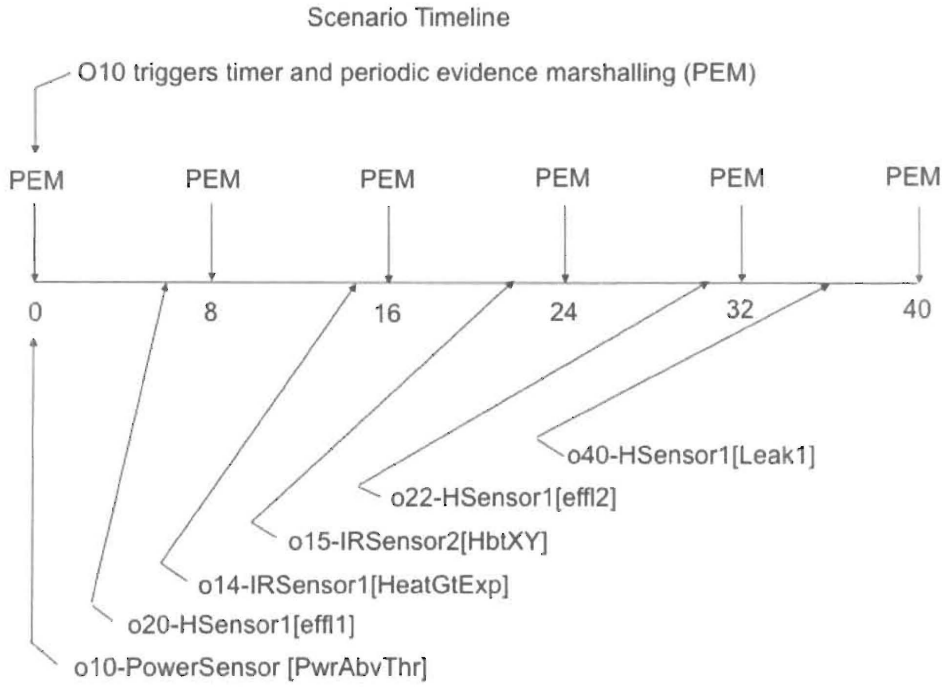

Figure 6. Scenario Timeline.

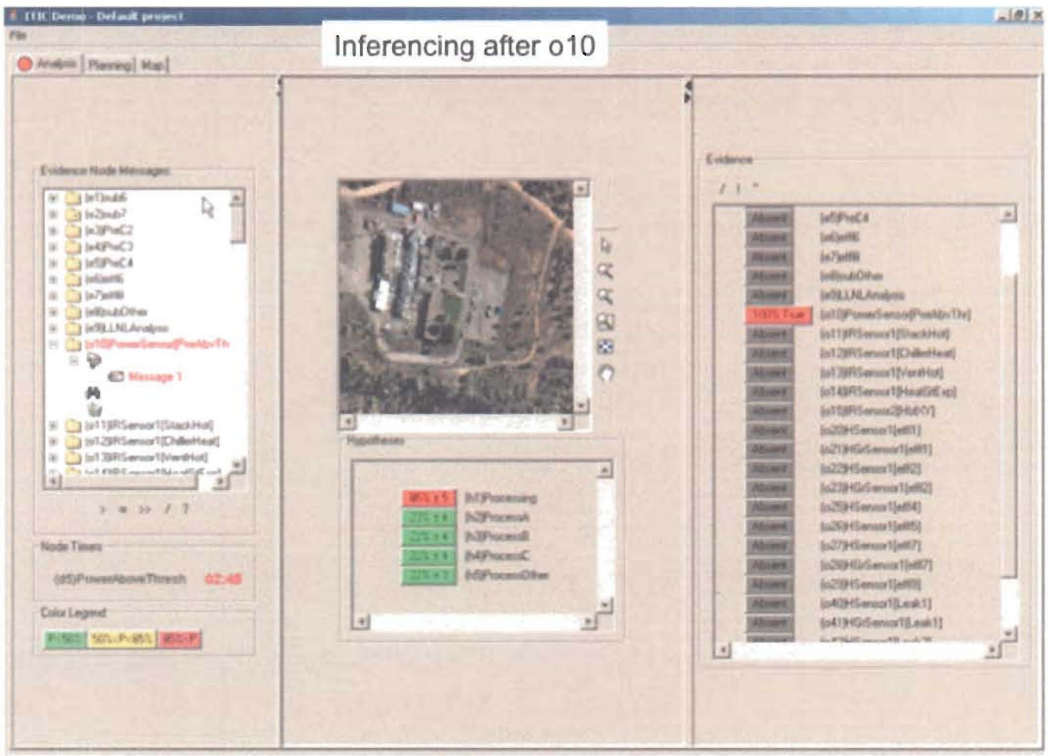

Figure 7. IKE Screen after first observation.

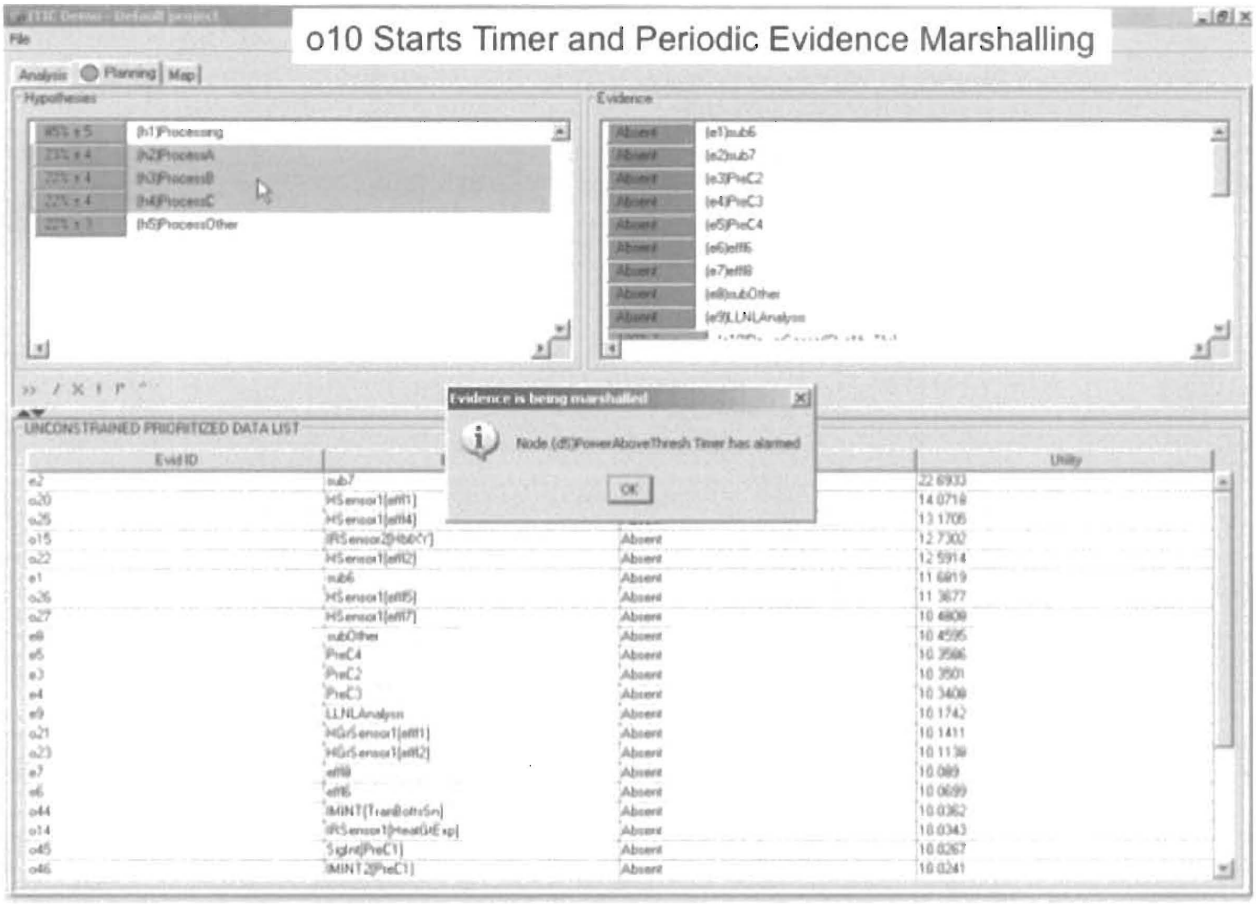

Figure 8. Evidence Marshalling.

Every time a new piece of evidence is collected, we then can calculate the "next best piece of evidence to collect," as shown in Figure 8. The uppermost pieces would be best to collect, as shown by our "utility" calculation on the far right. This is when one might be able to task, if different options are available.

The final results might be shown as indicated in Figure 9, depending upon the evidence received.

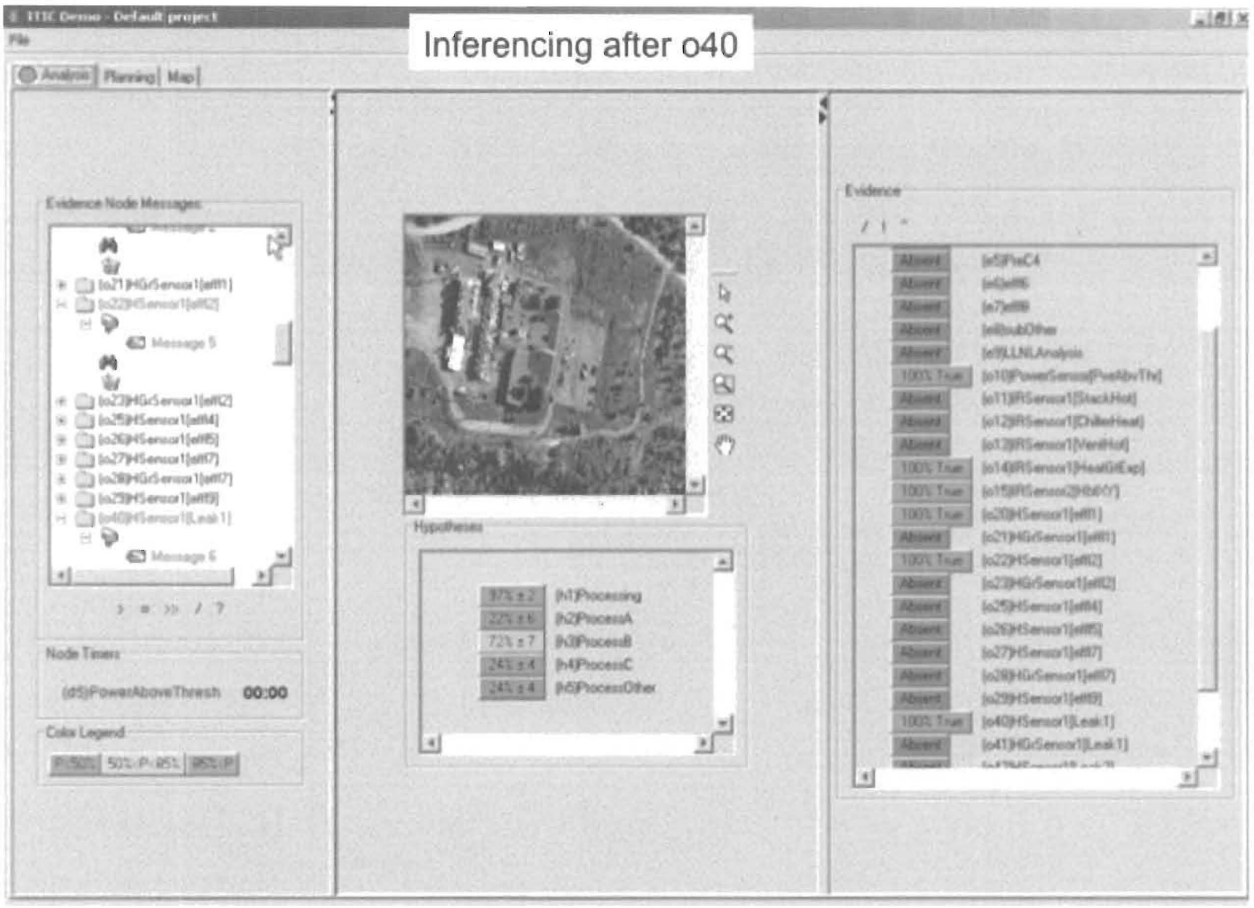

Figure 9. State of hypotheses after evidence collection.

5.0 Real-Time Monitoring

The second example (Figure 10 below) shows IKE embedded in a near-real-time architecture called the Remote Perimeter Surveillance (RPS) system used to monitor vehicular and pedestrian traffic on a remote canyon road. The RPS system integrates data from a distributed sensor network (DSN) consisting of seismic and acoustic nodes, with information from still cameras, radiation sensors, and from a video system that does motion detection and object tracking. The evidence generators consume these data streams to produce IKE evidence messages to set evidence into a BBN whose purpose is to autonomously determine the probability that the moving object is a vehicle (car, humvee, truck), a pedestrian, or some other entity such as an animal, and whether it is suspicious.

The upper left shows the video view from the canyon rim with a bounding box around the detected moving object (a car), middle left is a plot of the raw DSN data, and lower left shows a close-up view from a DSN still camera down in the canyon bottom. The upper right shows the DSN sensor nodes on a geographic situation display, with the nodes reporting shown in red, middle right shows raw/averaged DSN and video data, and bottom right shows the two main IKE hypothesis nodes and some evidence nodes for manual entry.

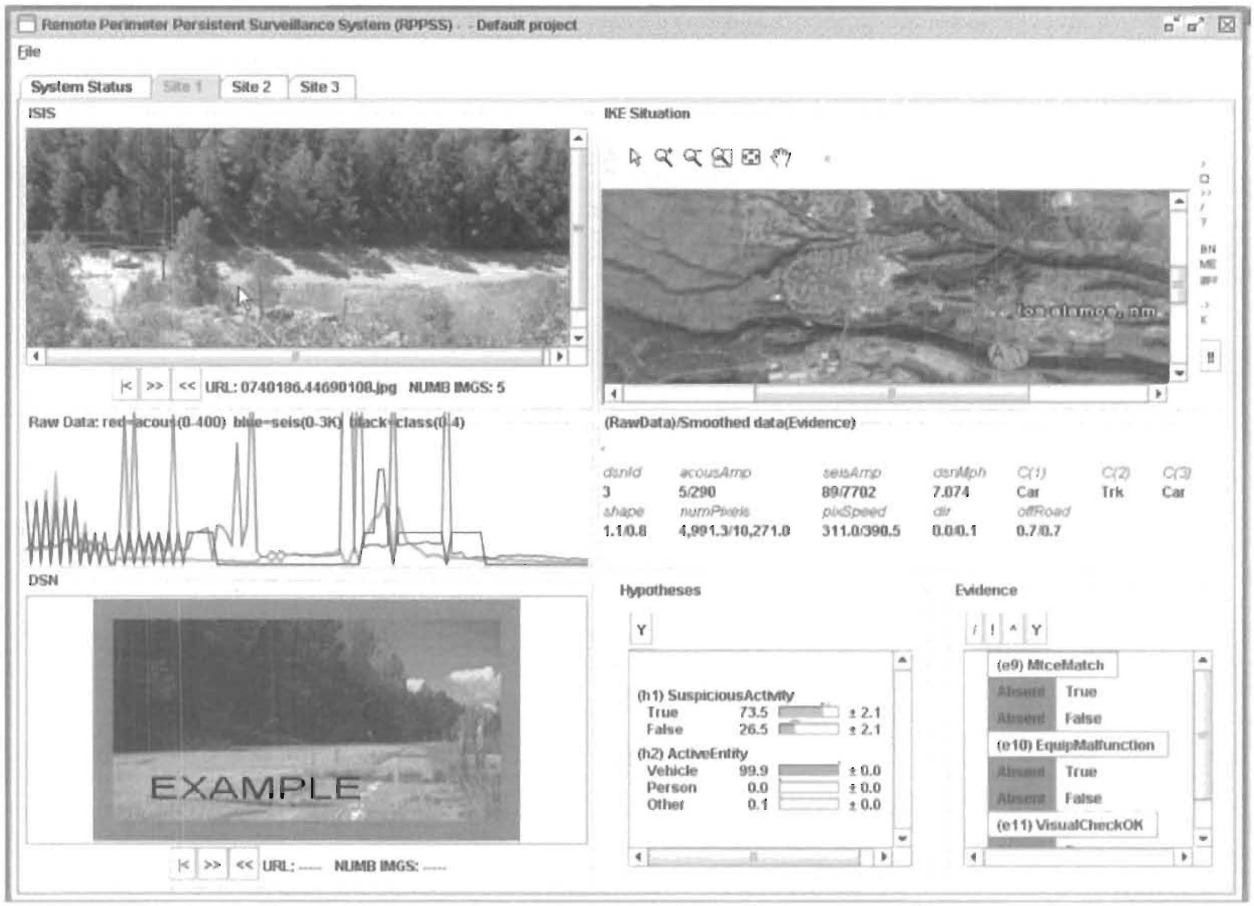

Figure 10. Remote Perimeter Surveillance System.

Even when a near-real-time system such as the RPS is running autonomously, the system can automatically perform evidence marshalling to guide additional evidence collection by (for example) triggering a special sensor to collect when the standard sensor suite is producing ambiguous results. Typically, the autonomous monitoring will alert an operator when something suspicious is detected, who may then choose to marshal and trigger additional evidence collection for the Bayesian analysis, or manually resolve any remaining ambiguity by doing an instant replay of imagery or data (Leishman and others, 2007).

6.0 Summary

The Integrated Knowledge Engine extends traditional Bayesian analysis by giving the modeler a way to express model parameter uncertainty at modeling time and by giving the analyst a way to express evidence uncertainty at run time. A Monte Carlo simulation wrapped around a traditional Bayesian analysis tool then allows these effects to be included in the results of the Bayesian inferencing. IKE also provides enhancements to optimize the collection of evidence and to understand when enough evidence has been obtained for a solid decision, and if not, a better understanding of the alternatives. Each of these enhancements provides a needed mechanism for many situations, specifically to arrive at a solution as quickly as possible with the least uncertainty.

We have demonstrated two examples of the utilization of IKE. We have worked on a variety of other problems, such as nuclear forensics, nuclear proliferation, IEDs, and other important scenarios.

7.0 Acknowledgements

This work has been supported by the Department of Energy, Department of Defense, and various intelligence agencies.

8.0 References

Gibson, W., Leishman, D. and E. Van Eeckhout, Jan 2009, "Applying Bayesian Belief Networks in Rapid Response Situations," Proceedings, 42nd Hawaii International Conference on System Sciences, LAUR-08-03952, 8 p.

Huang, C. and A. Darwiche, 1996, "Inference in belief networks: A procedural guide", International Journal of Approximate Reasoning, 15,(3), pp. 225-263.

Izraelevitz, D., Leishman, D., Martz, H. and W. Gibson, 2007, On Representing Second-Order Uncertainty in Multi-State Systems via Moments of Mixtures of Dirichlet, Accepted to the *Journal of Statistical Computation and Simulation*.

Jensen, F. and T. Nielsen, 2007, Bayesian Networks and Decision Graphs, Springer, 2nd ed., 448 p.

Leishman, D., Gibson, W., Rosten, E., Cai, M. May 2007, Remote Perimeter Persistent Surveillance System Screen Cast, Los Alamos National Laboratory, LAUR-07-6983.

Pearl, J., 2000, Causality: Models, Reasoning, and Inference, Cambridge University Press. 379 p.