

INVESTIGATING THE USE OF BAYESIAN NETWORKS TO PROVIDE DECISION SUPPORT TO MILITARY INTELLIGENCE ANALYSTS

Ken R. McNaught and Venkat V.S.S. Sastry
Cranfield University
Engineering Systems Department
RMCS Shrivenham, Swindon SN6 8LA, UK.
K.R.McNaught@cranfield.ac.uk

Bernard Ng
Singapore Armed Forces
MINDEF Building
Gombak Drive, Singapore 669645.

KEYWORDS: intelligence, military, probabilistic, Bayesian network, decision support, information fusion

ABSTRACT

In this paper we consider a typical military scenario where the intention of an enemy force is unknown, but there are a number of plausible hypotheses. As time passes, information in the form of various sightings and reports become available. We employ a Bayesian network, a type of probabilistic graphical model, to process these reports and update the probabilities of the various hypotheses in the light of the latest information. We also demonstrate the beneficial effect of incorporating 'negative' or false evidence on the plausibility of the various hypotheses.

INTRODUCTION

Despite the ever-increasing sophistication of combat simulations, the representation of various aspects of C4I and ISTAR (Intelligence, Surveillance, Target Acquisition and Reconnaissance) remains limited. In this paper, we examine the potential application of Bayesian networks to act as processors of intelligence reports, and provide a simplified, illustrative example. Such processors could be considered to be expert systems or intelligent agents, providing decision support to intelligence analysts within a Headquarters cell. Equally, they could be embedded within a combat simulation, representing one of the roles normally conducted by a military Headquarters' intelligence cell.

The network envisaged does not directly process low-level sensor and signal data, but operates on higher-level information provided by an intelligence analyst. It is intended to assist such analysts in making sense of the 'bigger picture', e.g. 'What is the enemy's most likely course of action given all of the indicators received and the background to the situation?' According to the JDL model, well-known within the data fusion community (see, for example, Llinas et al, 2004), such questions are associated with Level 3 in the hierarchy of data and information fusion. Answering such questions can

involve some very complicated reasoning, largely of a diagnostic nature. Unfortunately, many psychological studies (e.g. see Kahneman, Slovic and Tversky, 1982) have demonstrated that, through no fault of their own, humans often perform poorly at such tasks. This is one reason why there is such interest in computerised systems to assist with reasoning in complicated, uncertain domains such as medicine. Such systems are not intended to replace expert assessments, but to provide a second opinion. If that second opinion concurs with the expert's, then it will provide additional confidence in the assessment. On the other hand, if there is a strong disagreement, it may encourage the expert to re-examine the situation and check that nothing crucial has been overlooked.

Bayesian networks provide a powerful method of reasoning in domains where uncertainty is prevalent. As each new piece of evidence is received, the network can propagate its effects to whichever other nodes in the network are affected by it. This results in an updated set of beliefs regarding the key unknown variables of interest - usually one or more main hypotheses which are not directly observable until it is too late. Some other fields of application include medical diagnosis (e.g. Nikovski 2000), intelligent troubleshooting systems (e.g. Breese and Heckerman 1999) and data mining (e.g. Heckermann 1997).

BAYESIAN NETWORKS

A Bayesian network (BN) consists of a directed acyclic graph (DAG) and a set of conditional probability distributions for each node in the network. The graph comprises a set of nodes, with each node representing a proposition or variable within the domain of interest, and a set of directed arcs representing direct probabilistic dependencies between the variables. The absence of an arc between two variables is interpreted as a statement of conditional independence, i.e. the two variables are independent given some subset of the other variables in the network. For each variable without parents, we need to provide a prior probability distribution. For each variable with parents, we need to specify a conditional

probability distribution given each possible combination of parent states.

The conditional probability distributions which accompany a particular ordering of the variables in the graph, provide a compact way to specify the joint probability distribution over the entire set of variables, $U = \{A_1, A_2, \dots, A_n\}$:

$P(U) = \prod_i P(A_i | pa(A_i))$, where $pa(A_i)$ refers to the parents of variable A_i in the graph.

There are many potential orderings of variables in a network, and the ordering chosen for a BN should represent the assumed dependencies and independencies as efficiently as possible. This usually means that the direction of an arc should follow the direction of causality when the relationship between two variables is causal. So, it is the activities (or intent indicators) undertaken by the Red side which cause reports to be generated, the reports do not cause the activities to take place. Not all relationships in a BN have to be causal - weaker probabilistic dependencies will often be present. Exactly how such relationships should be represented and which way the arcs should be directed usually becomes clearer once the modeller has thought through their dependency implications. An invaluable guide in this respect is the d-separation criterion. See Pearl (1988) or Jensen (2002) for more details of this and for an introduction to Bayesian networks, more generally.

AN ILLUSTRATIVE SCENARIO

In this paper, we consider a simplified, general scenario in which the Blue force HQ is trying to ascertain the Red force's course of action.

Four possible courses of action (COAs) are considered in the scenario - main attack (M), advance (A), defend (D) and withdraw (W). It is assumed here that the Red force will only pursue a single course of action at any given time, although this assumption could be relaxed if it were thought necessary. This does not mean, however, that we do not permit the Red force to operate a deception plan.

Indicators of Enemy Intent

The intelligence staff tasked with inferring the Red force's most likely course of action, will have a number of cues, or key pieces of information, which they are interested in observing. These cues will be indicative of the course of action being taken by the Red force. Some of these indicators of enemy intent will be associated with a single course of action, while others will be

associated with more than one. Even when an indicator is associated with more than one COA, however, it might still provide greater support for one belief than another. The 20 indicators of enemy intent (IEI) considered here include the following:

- Increase in recce activities
- Increase in counter-recce activities
- Forward movement of supplies
- Establishment of airfields
- Forward movement of missiles
- Radio silence
- Use of smoke
- Erection of obstacles
- Preparation of dummy positions
- Increased anti-tank assets with forward units
- Evacuation of some services
- Destruction of various facilities

A BAYESIAN NETWORK MODEL OF THE SCENARIO

A basic Bayesian network for this scenario is shown in Figure 1. It consists of a hypothesis node, 'Enemy Intent', a layer of intermediate nodes representing a range of enemy intent indicators, and a layer of nodes representing battlefield intelligence reports relating directly to these indicators. A more complete network might also include various environmental variables.

The final number attached to each report node identifies which of the 20 enemy indicators it relates to. For example, node 'S3MA1' is a report from a Blue sub-unit that Red has increased its air and ground reconnaissance (recce) activities. This relates to the first enemy intent indicator in the middle layer of nodes. Note that setting the state of a report node to 'True' is not the same as setting the state of the corresponding enemy intent indicator to 'True'. Since incorrect and incomplete reports are to be expected, it is important to distinguish between the actual state of a variable and its perceived state. The conditional probability distribution of the report variable given the actual indicator variable will determine how much our belief in the indicator variable changes given a particular report.

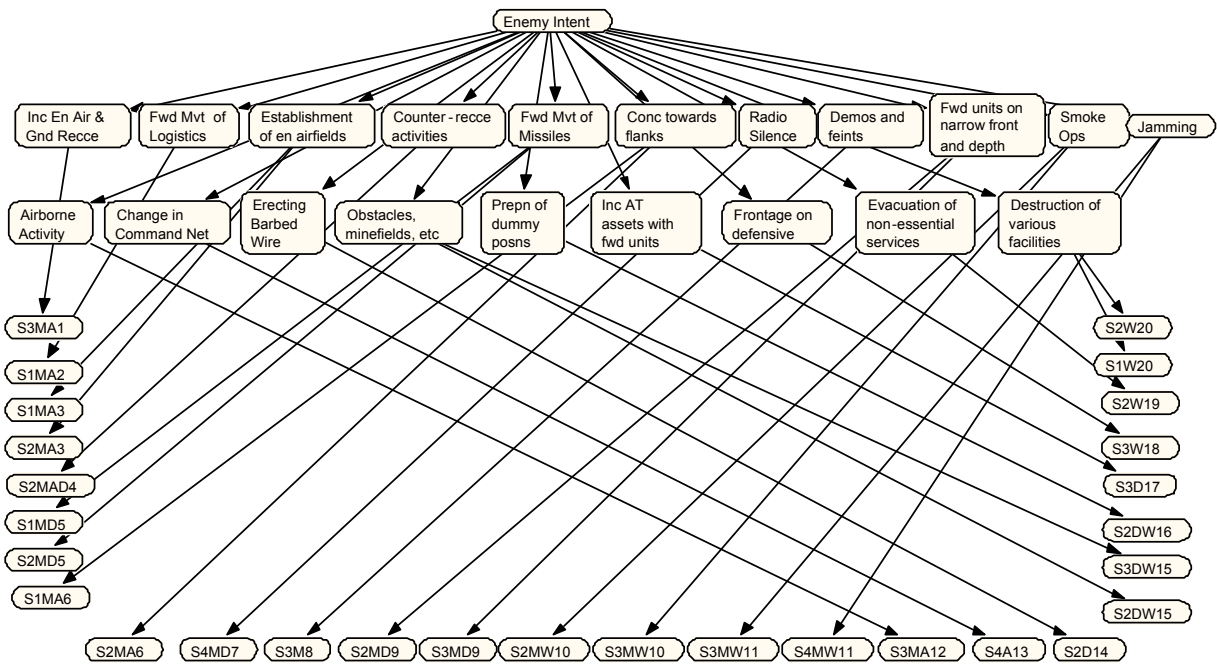


Figure 1: A Bayesian Network of the Relationships Between Enemy Intent Indicators and Intelligence Reports

While space constraints preclude a listing of the illustrative probability distributions contained within this model, a uniform prior distribution was assumed for 'Enemy Intent'. Obviously, the prior chosen reflects initial conditions and the knowledge of the Blue side. Then, as battlefield intelligence reports become

available, so the corresponding nodes will be instantiated. This leads to the probability distributions associated with the enemy intent indicator nodes and the hypothesis node being updated accordingly, as demonstrated in Figure 2.

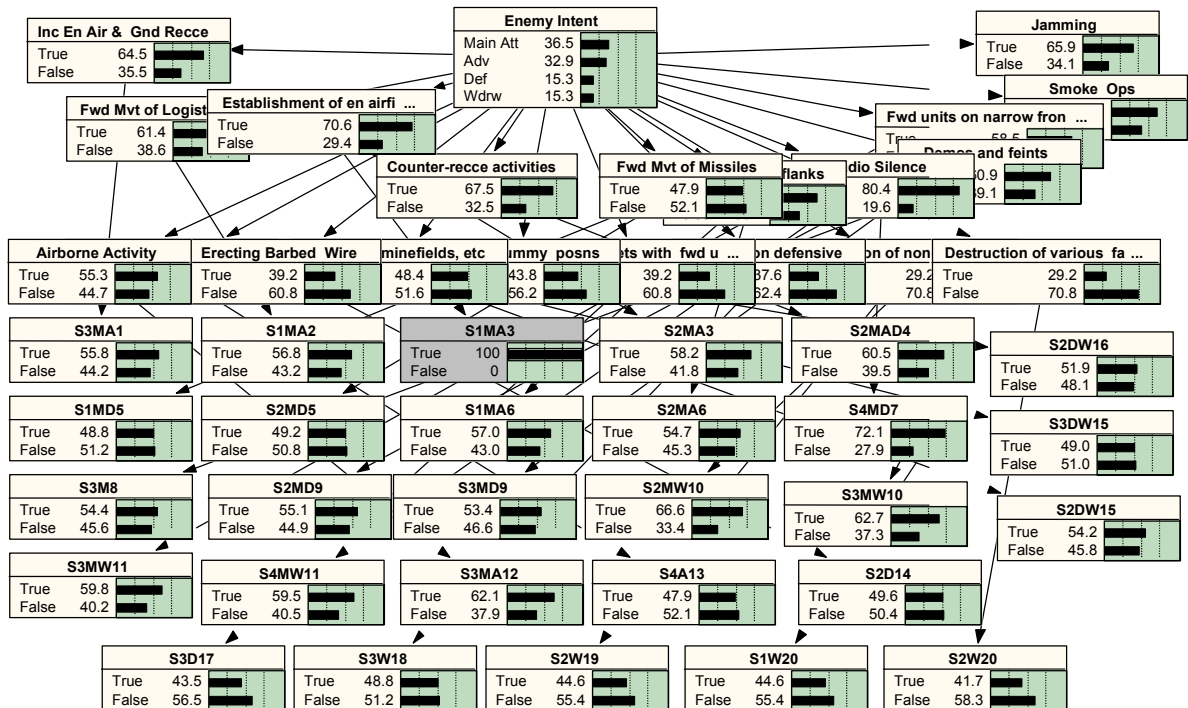


Figure 2: BN Showing Marginal Distribution of All Variables After Node 'S1MA3' is Set to 'True'

As can be seen in Figure 2, for example, if there is an air reconnaissance report of establishment of enemy airfields, represented in the network by setting node 'S1MA2' to 'True', then, denoting that evidence by 'E': $P(M | E) = 0.365$; $P(A | E) = 0.329$; $P(D | E) = 0.153$ and $P(W | E) = 0.153$.

The marginal distribution of every variable which is d-connected to 'S1MA2' is also updated at the same time.

Results of the First Experiment

We consider two vignettes to illustrate the approach. In the first one, the Red force is preparing to advance, while in the second one, the Red force is actually preparing to withdraw but attempts to conceal its intentions and deceive the Blue force into believing that it is preparing an attack.

The timeline for the first vignette is as shown in Table 1. Here, time is measured discretely in a number of steps. This denotes a simple chronological ordering of the events, which is sufficiently accurate for our purposes. It is not implied that the time steps are all of equal size. Also, there is a delay between an event's occurrence and the subsequent detection and reporting of that event to Blue HQ. As space constraints make it impractical to show the updated network after each time step, the graph in Figure 3 shows how the probability distribution of 'Enemy Intent' changes with time. Although this graph is shown as connected to make it easier to identify the

various states, the probability updates only occur at the discrete time steps.

Table 1: Timeline for Vignette 1.

Time Step	Actions Taken by the Red Side and Indicators Detected by the Blue Side
1	Both sides deploy air and ground recce.
2	Red deploys airborne forces to establish aux airfield; Blue sub-unit reports sighting of Red recce (S3MA1).
3	Red establishes aux airfield for CAS; Blue sub-unit reports sighting of Red airborne forces (S3MA12).
4	Red establishes counter-recce to cover advance route; Blue air and ground recce report sightings of Red aux airfield (S1MA3 and S2MA3).
5	Red formation strengthens flanks; Red forward movement of supplies; Blue ground recce reports Red counter-recce activities.
6	Red command net changes for advance; Blue air and ground recce report sightings of Red's strengthened flanks (S1MA6 and S2MA6); Blue air recce reports sightings of Red's forward movement of supplies (S1MA2).
7	Red begins advance; Blue Signals report Red change of command net (S4A13).




Figure 3: Probability Distribution of 'Enemy Intent' vs Time for Vignette 1

It is clear from the graph in Figure 3 that the correct enemy intent is eventually inferred, the final distribution being:

$$P(M | \text{All evidence}) = 0.46; P(A | \text{All evidence}) = 0.53;$$

$$P(D | \text{All evidence}) = 0.01; P(W | \text{All evidence}) = 0.005.$$

However, for a long time 'Main Attack' was considered the most likely intention, with 'Advance' only overtaking it towards the very end. Such an outcome could only be considered a partial success for the network.

The timeline for the second vignette is as shown in Table 2. A graph showing how the probability distribution of 'Enemy Intent' changes with time

Table 2: Timeline for Vignette 2.

Time Step	Actions Taken by the Red Side and Indicators Detected by the Blue Side
1	Blue establishes air and ground recce.
2	Red deploys air and ground recce as deception; Red increases counter-recce activities as deception; Red establishes dummy airfields as deception.
3	Red establishes demolition on bridges; Blue sub-unit reports sighting of Red recce (S3MA1); Blue ground recce reports Red counter-recce activities (S2MAD4); Blue air recce reports sighting of Red aux airfields (S1MA3).
4	Red conducts feint attacks; Blue ground recce report sighting of Red aux airfield (S2MA3) and demolition on bridges (S2DW15); Blue sub-unit reports local attacks (S3M8).
5	Red evacuates non-essential services; Blue sub-unit reports sighting of demolition on bridges (S3DW15).
6	Red employs smoke and jamming and a defensive frontage; Blue ground recce reports sighting of Red evacuation of non-essential services (S2W19) and Red's use of smoke (S2MW10); Blue sub-unit reports Red's use of smoke (S3MW10) and jamming (S3MW11); Blue Signals report Red's jamming (S4MW11); Blue sub-unit reports Red's defensive frontage (S3W18).
7	Red begins systematic destruction of bridges and commences withdrawal; Blue air and ground recce report sightings of Red destruction of bridges (S1W20 and S2W20).

for the second vignette is shown in Figure 4. Again, the correct inference is eventually made - this time the final probability distribution of 'Enemy Intent' is:

$$P(M | \text{All evidence}) = 0.39; P(A | \text{All evidence}) = 0.06;$$

$$P(D | \text{All evidence}) = 0.04; P(W | \text{All evidence}) = 0.51.$$

However, in common with the first vignette, the true enemy intent only became apparent towards the end. For much of the time, 'Main Attack' seemed the likelier option.

Results of a Second Experiment Incorporating Negative Evidence

The intentions of the Red force in these two vignettes were clearly difficult for the Blue HQ to identify. While the correct intentions were eventually identified, these came fairly late. It could be argued, however, that not all of the available relevant information was fed into the Bayesian network. In particular, events associated with one Red intent or another which were not observed to occur were assumed unknown. What would the effect be if after a suitable period of time, such events were reported as definitely not having occurred? This is investigated in a second experiment. The same underlying events are generated as in the first experiment, and the same positive intelligence reports are received at the same times. The difference is that in addition to the positive intelligence reports, there are now a number of 'negative' intelligence reports indicating that certain things have not been reported.

In deciding when to instantiate a report node with negative evidence, we have looked at the latest time we would expect a positive report to be received across the four possible states of Enemy Intent. If it has not been received by that time, we have instantiated a negative report for that indicator in the next time-step.

The revised results for the two vignettes, incorporating the effects of negative evidence, are shown in Figures 5 and 6.




Figure 4: Probability Distribution of Enemy Intent vs Time for Vignette 2




Figure 5: Probability Distribution of Enemy Intent vs Time for Vignette 1 Including Negative Evidence

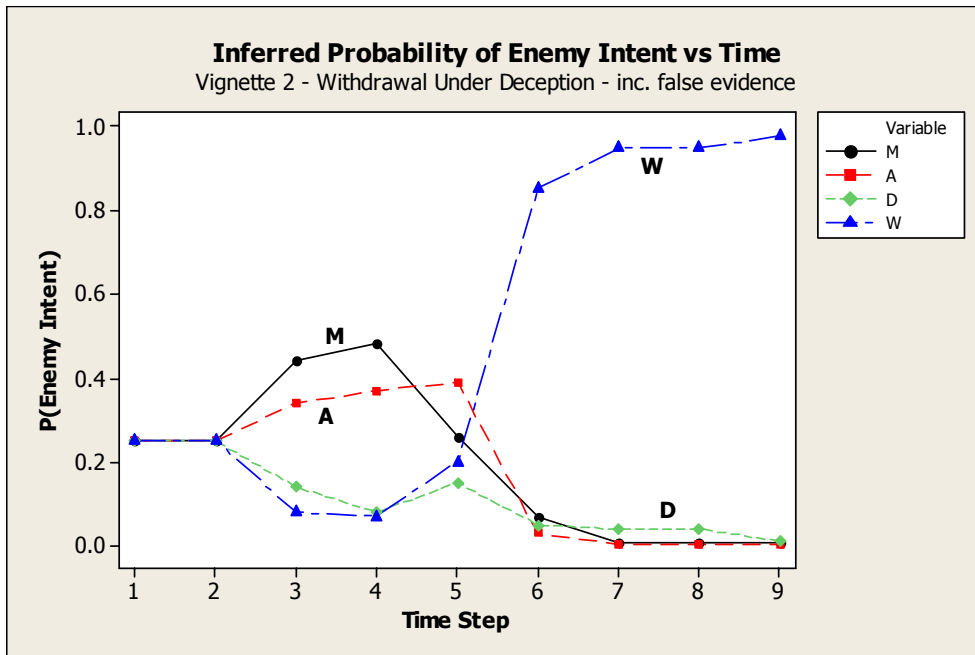


Figure 6: Probability Distribution of Enemy Intent vs Time for Vignette 2 Including Negative Evidence

Clearly, in both cases, the network performs much better when the negative evidence is also taken into account. Firstly, the final distribution of 'Enemy Intent' is more decisive in each case. In Figure 5, the final distribution of enemy intent is now given by: $P(M | \text{All evidence}) = 0.005$; $P(A | \text{All evidence}) = 0.994$; $P(D | \text{All evidence}) = 0.001$; $P(W | \text{All evidence}) = 0$. Similarly, in Figure 6, the final distribution of enemy intent is now given by: $P(M | \text{All evidence}) = 0.01$; $P(A | \text{All evidence}) = 0.005$; $P(D | \text{All evidence}) = 0.01$; $P(W | \text{All evidence}) = 0.975$. Secondly, the correct option is identified earlier by the network in both cases. While it is difficult to quantify the benefit obtained by identifying the true enemy course of action sooner, this could be addressed in a simulation study.

CONCLUSIONS

Using only positive evidence, the network sometimes has difficulty in discriminating between some of the alternatives. Although in the examples considered here, it eventually 'got it right', this was often very late. The performance was much improved when false evidence, indicating that certain indicators of enemy intent had not been observed, was also employed in the network. False findings at an observation node were only instantiated after the latest time they would normally have been

expected to have been observed had the indicator been present.

Further work will consider how timing information can be better exploited to avoid a sudden rush of false findings towards the end of a scenario causing large, discontinuous jumps in the probability distributions. The use of distributional information on detection times should permit smoother, continual updates in the

distributions over time, with discontinuous jumps occurring only when definite findings are observed. We will also attempt to quantify the benefits which such a decision support system can bring in terms of improved responsiveness. Simulation currently appears to be the most likely method of achieving this goal.

REFERENCES

- Breese, J. S. and Heckerman, D. 1999. "Decision-Theoretic Troubleshooting." *IEEE Transactions on Systems, Man & Cybernetics, Part A (Systems & Humans)* 26, No. 6, 838-842.
- Heckerman, D. 1997. "Bayesian Networks for Data Mining." *Data Mining and Knowledge Discovery* 1, No. 1, 79-120.
- Jensen, F.V. 2002. *Bayesian Networks and Decision Graphs*. Springer-Verlag, New York.

Kahneman, D.; Slovic, P.; and Tversky, A. (Eds.) 1982. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge University Press, Cambridge, Mass.

Llinas, J.; Bowman, C.; Rogova, G.; Steinberg, A.; Waltz, E.; and White, F. 2004. "Revisiting the JDL Data Fusion Model II." In *Proceedings of the 7th International Conference on Information Fusion*, Stockholm, 1218-1230.

Nikovski, D. 2000. "Constructing Bayesian Networks for Medical Diagnosis from Incomplete and Partially Correct Statistics." *IEEE Transactions on Knowledge and Data Engineering* 12, No. 4, 509-516.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, San Mateo, California.

AUTHOR BIOGRAPHIES

Ken R. McNaught is a lecturer in Operational Research (O.R.) at Cranfield University's College of Defence Technology in Shrivenham, UK. He has an MSc in O.R. from Strathclyde University and a PhD in O.R. from Cranfield University. His current research interests include simulation, combat modelling and decision support, particularly making use of probabilistic graphical approaches such as Bayesian networks and influence diagrams.

Bernard Ng holds the rank of Major in the Singapore Armed Forces, in which he has served for 12 years. He holds a BSc in Computer Science from the National University of Singapore and an MSc in Defence Technology from Cranfield University. He is currently operating as a staff officer SO1 in the area of C3I.

Venkat V. S. S. Sastry obtained his PhD (1980) in Applied Mathematics from IISc Bangalore, India. Venkat worked on analytical models of propagation of sound over ground, in underwater sediments (1981-86) as a post doctoral research fellow. He joined Cranfield University at the Royal Military College of Science, Shrivenham in 1989 as a lecturer. He is currently the Director of Applied Mathematics and Scientific Computing. His current interests include scientific visualization, applications of Intelligent Systems and Virtual Environments for Training, gesture interpretation in virtual environments for training applications and application of agents for decision support.