

Bayesian Network meta-models from combat simulation for Defence decision analysis

Abida Shahzad,^a Steven Mascaro,^b Thang Cao^c and Kevin B. Korb^a

^aFaculty of IT, Monash University, VIC, Australia, 3800

^bBayesian Intelligence Pty Ltd, Melbourne, Australia

^cDefence Science and Technology Group, Australia

Email: abida.shahzad@monash.edu

Abstract: Defence Science and Technology Group (DST) is investigating the characteristics of Land Combat Vehicles (LCVs) with a view to improving purchasing decisions between different LCV options. Using a combination of closed-loop combat simulation (i.e., one run without any human interventions) and subject matter experts (SMEs), these studies seek to understand the close combat capability factors for a future LCV. These factors are being investigated for different kinds of physical environment and different levels of enemy lethality, with each LCV's lethality, survivability, signature (the ease with which the enemy identifies friendly forces) and knowledge acquisition (friendly forces identifying enemies) being measured, yielding a multi-dimensional view of the different LCVs and their performance. Our study describes a new approach, developing a Bayesian network (BN) meta-model (i.e., a model based upon the combat simulation model; see Fig 1) that combines those multiple dimensions in a single multicriterial decision model.

Both deterministic and stochastic methods have been developed for multicriteria decision analysis (MCDA). Limitations to prior approaches include treating utilities as independent of each other and also very limited ability to deal with uncertainty about preferences. Bayesian networks offer advantages in these regards, so here we develop an alternative approach using a BN meta-model for decision making and evaluating the operational impact of different LCV choices. BNs support both the prediction of operational results and their causal explanation. These BNs were developed both from expert elicitation and causal discovery (data mining) from combat simulation data, and enable causal visualization, sensitivity analysis and optimization. We evaluate the performance of some of these BNs against alternatives, using common accuracy-based evaluation metrics. Please note that the example and data presented in this paper are not representative of the real LCV options and testing scenario, they are used for methodology illustration only.




Figure 1. A Bayesian network meta-model (from the PC algorithm).

Keywords: Bayesian network modelling, multicriteria decision analysis, combat simulation and model evaluation

1 INTRODUCTION

In multicriteria decision analysis (MCDA), decision makers (DMs) choose from m alternatives using n criteria. Often criteria are themselves composed of weighted sub-criteria. The preferences (weights, utilities) associated with the (sub)criteria may be determined by interviews, surveys, etc. with SMEs, direct measurements, or simulations. Whatever their sources, real world applications must cope with uncertain and incomplete data. The weights are also frequently uncertain and occasionally incomplete. In our study, the criteria utilities and weights are evaluated using both combat simulation and SMEs.

Deterministic methods of ranking and selecting the “best option” from simulation has been extensively studied by Nelson et al. (2001) for a single response metric, and there are some studies, e.g., Cao et al. (2010) and Emond (2006), which apply different MCDA methods to multicriterial problems. AHP weighting is the most common deterministic weighting approach in MCDA defence problems (Cao et al., 2010; Nguyen and Cao, 2017) due to its capacity to establish the weights of attributes in a systematic and robust manner and allow the DMs to check the consistency and consensus between the rankings of the criteria in group decision making.

However, DST deals extensively with MCDA problems that require handling uncertainty about both attributes and utilities. Several other MCDA methods do support handling uncertain or inaccurate criterial information. For example, ELECTRE (Roy, 1996) and PROMETHEE (Brans et al., 1984) handle different kinds of uncertainty using threshold models. In multi-attribute utility theory (MAUT) the uncertainty is represented using probability distributions, and alternatives are evaluated based on their expected utility (Keeney and Raiffa, 1993). Of these, the MAUT-based methods are able to handle criteria uncertainties, but not preference uncertainties. As in MAUT, the Stochastic Multicriteria Acceptability Analysis (SMAA) methods (see e.g., Lahdelma et al., 1998, and Lahdelma and Salminen, 2001) represent uncertain or inaccurate criteria through probability distributions and are thus able to handle arbitrary dependencies between criteria uncertainties.

We here present an alternative approach using BN meta-models for decision making and evaluate their performance against alternatives, using common accuracy-based evaluation metrics.

2 COMBAT SIMULATION STUDY

The combat simulation used in these studies is based on a mission requiring the rapid seizure of an urban fringe area. The Blue force (the side for which LCV options were being assessed) was a Combat Team (CT) of three infantry platoons, a Tank troop and two direct fire support vehicles. The Red force was one third the size of Blue and its mission was to delay Blue by 24 hours. The CT assembly area was 10km from the objective in a heavily vegetated area, with one viable road to the objective. Red planned an ambush along the road.

Three LCV options were evaluated: (A) Wheeled, Weapon E, Armour X; (B) Tracked, Weapon F, Armour Y; (C) Tracked, Weapon E, Armour Y. Armour protection was measured by comparing casualty rates and kinds of infantry and vehicle damage; lethality measures were the same but applied to enemy platforms and infantry. The signature metrics reflected the total number of Blue force entities detected and their range at detection. Knowledge acquisition captured the same information in relation to detecting Red. A binary mission success variable recorded whether Blue achieved its mission, the seizure of its objective within 24 hours before reaching attrition thresholds. SMEs were consulted in order to generate appropriate plans and tactics for both sides. These plans were then encoded in the COMBATXXI simulation, a stochastic, entity-based closed-loop combat simulation developed by the US Army (Balogh and Harless, 2003). The simulation was repeated 10,000 times for each LCV. The output variables measured for each side were: LCVs




Figure 2. Simulation results (means and confidence intervals). The sample size is 10,000, so intervals are very small.

Killed; Tanks Killed; Infantry Killed; Detections of entities; Blue Mission Success. The mean and confidence interval of the metrics are shown in Figure 2.

3 USING BAYESIAN NETWORK META-MODELS

We developed Bayesian networks (BNs) to model combat effectiveness, as they allow qualitative and quantitative representations of the domain and specifically support decision analysis (Cao and Chau, 2015). A Bayesian network is a probabilistic graphical model, namely a directed acyclic graph (DAG), $G = \{X, E\}$, with $X = \{X_1, \dots, X_n\}$ a set of n random variables, each with a conditional probability table (CPT) representing conditional dependencies of the node X_i on its parents pa_i and $E = \{e_1, \dots, e_m\}$ is a set of m edges representing those direct dependencies.

Our chosen BN tool, GeNIe, provides a user friendly interface for developing BN models. It also includes the SMILE Engine (Druzdzel, 1999) for programmatic access to Bayesian inference with the developed BN model. We discretized the simulation data prior to applying two of the causal discovery algorithms within GeNIe, PC (Spirtes et al., 1993) and a Bayesian search based on K2 (Herskovits and Cooper, 1992) that uses BDeu (Bayesian Dirichlet equivalent uniform) as the metric (Heckerman and Geiger, 1995).

3.1 PC ALGORITHM

PC is a constraint-based algorithm that uses classical significance testing for dependencies to identify local structural constraints. It is widely used for causal discovery because it is easily explained and runs quickly. It was developed by Spirtes et al. (1993). PC starts with a fully connected *skeleton* network, where every node is connected to every other, removing edges according to any significance tests reporting conditional independence. Figure 1 shows the network learned using the PC algorithm.

3.2 K2

This structure learning algorithm is one of the earliest BN learning algorithms, developed by Herskovits and Cooper (1992). It uses greedy search in an iterative process to obtain the best DAG structure per a Bayesian metric. This algorithm requires a variable order before the search process can start, but is often combined with a minimum weighted spanning tree algorithm to provide an ordering.

4 META-MODEL VALIDATION

For validating meta-models we used two common measures directly supported by GeNIe, namely predictive accuracy and the log score function, applying them in *five-fold cross validation*.

	5 Folds					Accuracy	Base rate	Difference	Log Score
	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5				
Mission Success	0.9755	0.9645	0.9666	0.9643	0.9545	0.96508	0.893	0.07	-0.108
Blue LCV Killed	0.6606	0.6608	0.6543	0.6648	0.6683	0.66176	0.512	0.15	-0.673
Blue Tank Killed	0.6116	0.599	0.6068	0.6205	0.594	0.60638	0.554	0.052	-0.694
Blue Infantry Killed	0.6268	0.6176	0.6176	0.625	0.6245	0.6223	0.555	0.067	-0.806
Red LCV Killed	0.6665	0.6725	0.662	0.6648	0.682	0.66956	0.675	0	-0.659
Red Tank Killed	0.7586	0.7538	0.7406	0.748	0.7591	0.75202	0.729	0.023	-0.711
Red Infantry Killed	0.5943	0.5825	0.583	0.5676	0.58	0.58148	0.487	0.094	-0.841
Red Detects Blue	0.7065	0.6923	0.6953	0.7086	0.6936	0.69926	0.667	0.0323	-0.63
Blue Detects Red	0.972	0.9543	0.9698	0.9693	0.972	0.96748	0.942	0.025	-0.083

Figure 3. Predictive Accuracy for the PC meta-model. Bold indicates a statistically significant improvement in the Overall result compared to the Base Rate.

Figure 3 shows the accuracy results for each node in the model generated using the PC algorithm (i.e., the model in Figure 1). The primary output node of interest, *Mission Success*, reports a mean accuracy across all folds of about 0.965, which is statistically significantly higher than the Base Rate of 0.893 (i.e., the overall *Mission Success* rate across the simulations). The difference between the two is the “effect size” — the improvement obtained in predictive accuracy when using the observed values of the other variables to predict

	5 Folds					Accuracy	Base rate	Difference	Log Score
	Fold 1	Fold 2	Fold 3	Fold 4	Fold 5				
Mission Success	0.955	0.956	0.971	0.9696	0.954	0.96112	0.904	0.06	-0.133
Blue LCV Killed	0.6503	0.633	0.639	0.644	0.6355	0.64036	0.517	0.123	-0.664
Blue Tank Killed	0.606	0.607	0.582	0.595	0.588	0.5956	0.551	0.045	-0.699
Blue Infantry Killed	0.6306	0.624	0.622	0.6275	0.629	0.62662	0.557	0.07	-0.793
Red LCV Killed	0.672	0.677	0.6721	0.668	0.685	0.67482	0.675	0	-0.647
Red Tank Killed	0.757	0.7535	0.746	0.747	0.7605	0.7528	0.739	0.014	-0.719
Red Infantry Killed	0.565	0.569	0.594	0.555	0.5375	0.5641	0.486	0.08	-0.87
Red Detects Blue	0.697	0.686	0.694	0.701	0.681	0.6918	0.668	0.02	-0.626
Blue Detects Red	0.9705	0.965	0.9583	0.9603	0.975	0.96582	0.954	0.012	-0.0802

Figure 4. Predictive Accuracy for the K2 meta-model. Bold indicates a statistically significant improvement in the Overall result compared to the Base Rate.

whether the mission will be successful. We also performed five-fold cross validation on the network generated by the K2 algorithm (not shown), resulting in Figure 4. The mean accuracy of about 0.96 for *Mission Success* is again statistically significant, but slightly worse than that using PC. In both models, the other variables show some improvement in predictive accuracy as well, although not in all cases; the log scores are closely aligned with predictive accuracies. Since both models are similar, we limit subsequent discussion to the PC model.

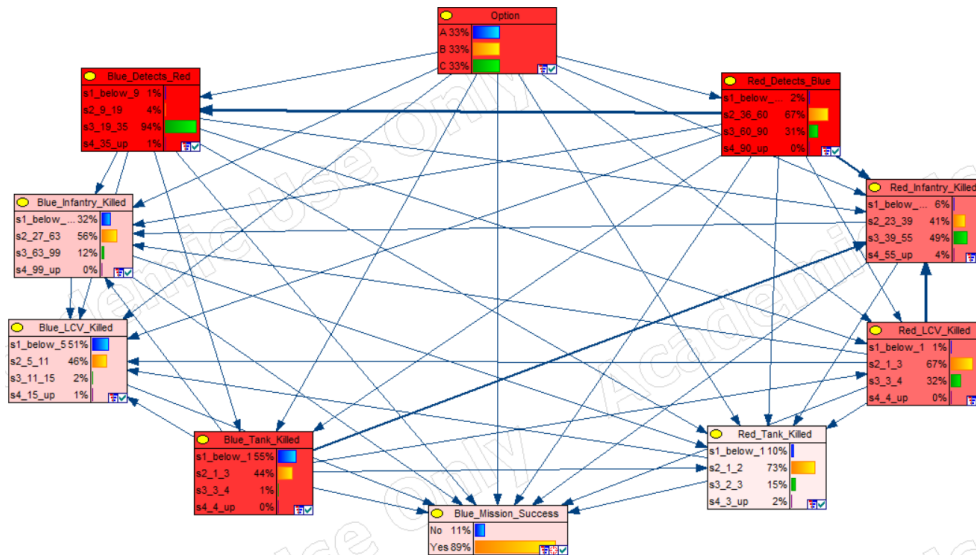


Figure 5. Sensitivity Analysis using causal visualization of the PC model.

5 CAUSAL VISUALIZATION

Figure 5 illustrates one kind of visual sensitivity analysis that can be performed using BNs. The degree of saturation of red indicates a node’s relative importance to *Mission Success*, with greater saturation being more important, while the thickness of the arrows showing the paths of greatest influence. These have been assessed using the mutual information between *Mission Success* and the other nodes. Examination of the visualization can readily reveal interesting points from the simulation. For example, it shows that detection of enemy forces, together with *Blue Tank Killed*, are the most important factors for *Mission Success*. Furthermore, Blue detecting Red is strongly related to Red detecting Blue. The model asserts that Red detecting Blue *causes* Blue detecting Red, however, that needs to be taken with a “grain of salt”: since that causal arrow is “covered” (i.e., the two detection nodes have a common parent), causal discovery algorithms cannot determine which way the arrow should go. Here, it can simply be read as implying a strong association between the two detection nodes.

(The causal discovery literature explains these issues in depth; see, e.g., Part II of Korb and Nicholson, 2010.)

6 WHAT-IF ANALYSIS

Causal BN models support arbitrary hypothetical or counterfactual analyses. That is, using them we can ask any kind of "What if" question we like. For example, what if we can protect all the Blue LCVs, Blue Tanks and Blue Infantry, what then will be the chance of Blue mission success?

If little to no harm comes to any of our Blue forces, i.e., *Blue LCV Killed*, *Blue Tank Killed* and *Blue Infantry Killed* are all in their lowest possible states, then we can see in Figure 6 that the success rate is about 99 percent. Note that this is strictly an *observational* analysis: it reveals a probabilistic dependency between observing Blue forces being protected and *Mission Success*, without revealing that *acting* specifically to protect Blue forces would lead to *Mission Success*. This is an example of correlation not being the same as causation.

A true causal BN can also answer causal hypotheticals: by manually altering Figure 6 by cutting all in-bound (parental) arrows to the nodes *Blue LCV Killed*, *Blue Tank Killed* and *Blue Infantry Killed* and then setting them to the states of interest (Blue forces protected), you will discover that directly intervening to protect Blue forces (as if by an "act of God") will indeed lead to greater *Mission Success* — increasing to about 96%, about 7% above the base rate. The observational and causal reasoning result in the same qualitative conclusion in this case, but in general observation does not substitute for causal intervention, any more than a Lecturer coloring her hair white would get her promoted to Professor. The ability to support both observational and causal reasoning is a particular advantage to using causal BNs for decision making that other approaches lack.

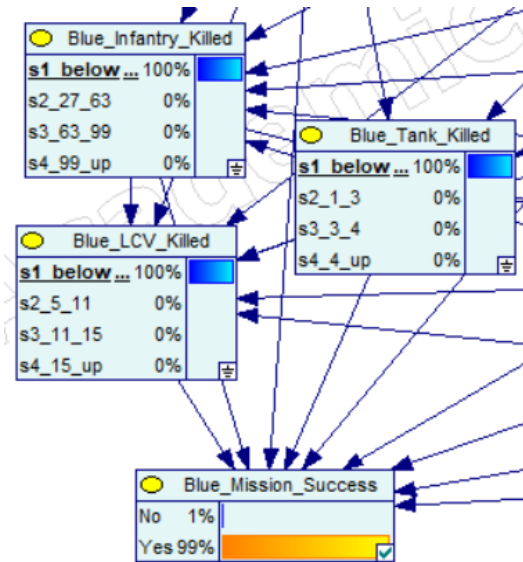


Figure 6. What-if Analysis

7 EXTENSION TO UTILITIES

Our BN can inform us about the chance of mission success given our three LCV options. If our sole concern was mission success, this would be sufficient but we are also concerned with other outcomes. We would like to minimize the loss of our own forces and equipment and inflict the same on the enemy. Each of these outcomes will matter to the DMs to a different degree and should therefore be weighted in our decision appropriately.

Figure 7 represents the model after adding utility nodes to each outcome of interest. There are six utilities associated with the forces (infantry, LCVs and tanks) of both Blue and Red (*BlueInfantry*, *BlueLCV*, *BlueTank*, *RedInfantry*, *RedLCV*, *RedTank*) and one associated with *MissionSuccess*. A final utility node is for estimating the effectiveness of the mission in a broader context, beyond just the current mission (*MissionEffectiveness*), which here is the weighted sum of all other utility nodes.

Each force utility node is a simple numerical translation of its outcome. For example, the utility node *RedInfantry* maps each state of *Red Infantry Killed* to the mid-point value of the state's range (so 0–23 is mapped to 11.5, 23–39 is mapped to 31, and so forth). The same applies to Blue force utility nodes, except that the associated utility is negative. *Mission success* is simply assigned 1 for success, and 0 for failure.

Of course, each such utility is on an independent scale; we need to combine them in the final *MissionEffectiveness* utility node. The combining weights, then, are similar to regression coefficients, but dependencies are being handled further up in the network. Since there is no obvious scale for the final utility node, we simply measured in terms of number of Blue infantry; e.g., -3 for mission effectiveness would be equal in value to the loss of 3 infantry. We describe these units as 'utiles'.




Figure 7. BN Meta Model with Utilities

The combining weights need to reflect the values of the DMs. The approach we used to estimate them was a mixture of the Trade-off Method (Schoemaker and Waid, 1982; Belton and Stewart, 2002) and Swing Weights (Edwards and Baron, 1994), where the weights were estimated via ratio values between the base criteria value and other criteria values. The weights were chosen after consultation over equipment capabilities: Blue Infantry 1; Blue LCV 15; Blue Tank 8; Red Infantry 0.2; Red LCV 7.5; Red Tank 4; Mission Success 100.

The critical issue of this decision problem is the trade-off between the competing objectives: reducing our force casualties, achieving mission success, protecting important assets such as LCV and tanks, and decreasing enemy capability. Per above, we took Blue infantry killed as the base unit, a utility of -1. We give the corresponding Red utility nodes lower weight. While there is an advantage to incurring enemy losses, they are less likely to be considered as important as losses to our own side beyond the current mission. For estimating the value of LCVs and tanks, we took into account their strategic value. LCVs are particularly valuable for missions of this class. Again, Blue losses are expected to be more significant to us than Red losses, but given the strategic value of LCVs, the discount for Red LCVs versus infantry is reduced. The same considerations were applied to tanks. Finally, mission success is likely to be highly context dependent — here, we have arbitrarily set it to equivalent to 100 Blue infantry. Please note the utility weight values presented here may not reflect Australian Defence strategic guidance; the relevant point is that in any concrete analysis, some weighting must be adopted.

With the weights in place, we can assess LCV options. The results (expected values) of *MissionEffectiveness* are: *A* -30.14; *B* 8.23; and *C* 9.81. Clearly *A* is ruled out. Options *B* and *C* are very close; any choice between them should be based upon a wider context, possibly requiring additional simulation scenarios.

8 CONCLUSION

In this paper, we explored three Land Combat Vehicle (LCV) options, evaluating them using multiple dimensions of utility, such as Blue force survivability and lethality. Bayesian Network (BN) meta-models were learned from combat simulation data by applying two common causal discovery algorithms. The resultant causal BN meta-models were then validated on the data; they could also be validated using expert opinion or a combination of expert opinion and data (see, e.g., Korb et al., 2013). Once you have a validated causal BN meta-model, you can explore it for insights into the domain using sensitivity analysis, including "what if" hypothetical analysis. Our causal BN meta-models were then further extended using expert-elicited utilities, yielding a causal decision network. We then used the causal decision network to determine the expected utility

ties for each LCV option, finding that for the simulated environment used in this study, option *A* had the worst expected utility, while options *B* and *C* were of about equal best value. To make a decision between *B* and *C* additional simulation environments will likely be necessary. This study shows the method and value of using BN meta-models to interpret combat simulations, or other simulations, with multiple dimensions of utility.

In the future, we intend to extend these methods by applying causal discovery algorithms to the data extended by elicited utilities. That is, experts' elicited utilities are themselves measurements of random variables (reporting their subjectively assessed utilities associated with aspects of each scenario). The variations and co-variations between utilities have a causal explanation as much as other variables, so causal discovery can provide insights into how and why they are related to each other, as well as to other random variables. We anticipate that such research may reveal new insights into multi-criteria decision making.

REFERENCES

- Balogh, I. and G. Harless (2003). An overview of the COMBAT XXI simulation model: A model for the analysis of land and amphibious warfare. In *Proceedings of 71st Military Operations Research Society Symposium*.
- Belton, V. and T. J. Stewart (2002). *Multiple Criteria Decision Analysis: An Integrated Approach*. Boston: Kluwer.
- Brans, J. P., B. Mareschal, and P. Vincke (1984). PROMETHEE: A new family of outranking methods in multicriteria analysis. *Operational Research* 84, 477–490.
- Cao, T. and W. Chau (2015). Bayesian network model for analysing operational impact of land combat vehicle. In *Proceedings of 12th Engineering Mathematics and Applications Conference (EMAC)*, Adelaide, Australia.
- Cao, T., A. Coutts, and B. Pietsch (June 2010). Multi-attributes Utility Theory and Statistical Analysis for Defence Future Vehicle Options. *ASOR Bulletin* 29(2), 69–81.
- Druzdzel, M. J. (1999). SMILE: Structural Modeling, Inference, and Learning Engine and GeNIe: A Development Environment for Graphical Decision-theoretic Models. In *Proceedings of 16th National Conference on Artificial Intelligence, AAAI*, Menlo Park, CA, USA, pp. 902–903.
- Edwards, W. and H. F. Baron (1994). SMARTS and SMARTER: Improved simple methods for multi-attribute utility measurement. In *Organizational Behaviour and Human Decision Processes*, Volume 60, pp. 306–325.
- Emond, E. J. (2006). *Developments in the Analysis of Rankings in Operational Research*, Volume 37. ORD Project Report (PR2000/13). Defence R&D Canada, Centre for Operational Research and Analysis. Directorate of Operational Research Corporate, Air & Maritime.
- Heckerman, D. and D. Geiger (1995). Learning Bayesian networks: a unification for discrete and Gaussian domains. In P. Besnard and Hanks (Eds.), *UAI 95*, San Francisco, pp. 274–84. Morgan Kaufman.
- Herskovits, E. and G. F. Cooper (1992). A Bayesian method for the induction of probabilistic networks from data. *Machine Learning* 9(4), 309–347.
- Keeney, R. L. and H. Raiffa (1993). *Decisions with multiple objectives; preference and value tradeoffs*. Cambridge University Press, Cambridge, UK.
- Korb, K., N. Geard, and A. Dorin (2013). A Bayesian approach to the validation of agent-based models. In A. Tolk (Ed.), *Ontology, epistemology, and teleology for modeling and simulation*, pp. 255–269. Springer, Berlin.
- Korb, K. B. and A. E. Nicholson (2010). *Bayesian Artificial Intelligence*. Boca Raton, FL: CRC Press.
- Lahdelma, R., J. Hokknen, and P. Salminen (1998). Stochastic multiobjective acceptability analysis. *European Journal of Operational Research* 106, 137–143.
- Lahdelma, R. and P. Salminen (2001). Stochastic multiobjective acceptability analysis for group decision making. *Operational Research* 3, 444–454.
- Nelson, B. L., J. Swann, D. Goldsman, and W. Song (2001). Simple procedures for selecting the best simulated system when the number of alternatives is large. *Operations Research* 49(6), 950–963.
- Nguyen, M. T. and T. Cao (2017). A hybrid decision making model for evaluating land combat vehicle system. In *MODSIM 22nd International Conference on Modelling and Simulation*, pp. 1399–1405.
- Roy, B. (1996). *Multicriteria Methodology for Decision Aiding*. Dordrecht, The Netherlands: Kluwer Academic Publishers.
- Schoemaker, P. J. and C. C. Waid (1982). An experimental comparison of different approaches to determining weights in additive value models. *Management Science* 28(2), 182–196.
- Spirtes, P., C. Glymour, and R. Scheines (1993). *Causation, Prediction, and Search*. Springer, New York.