Proceedings of the 2005 International Conference on Simulation and Modelling
V. Kachitvichyanukul, U. Purintrapiban, P. Utayopas, eds.

PARAMETERISING BAYESIAN NETWORKS: A CASE STUDY IN ECOLOGICAL RISK ASSESSMENT

Carmel A. Pollino

Water Studies Centre
School of Chemistry
Monash University
Clayton, 3800 AUSTRALIA
carmel.pollino@sci.monash.edu.au

Owen Woodberry, Ann Nicholson, Kevin Korb

School of Computer Science
and Software Engineering
Monash University
Clayton, 3800 AUSTRALIA
{owenw,annn,korb}@csse.monash.edu.au

ABSTRACT

Most documented Bayesian network (BN) applications have been built through knowledge elicitation from domain experts (DEs). The difficulties involved have led to growing interest in machine learning of BNs from data. There is a further need for combining what can be learned from the data with what can be elicited from DEs. In previous work, we proposed a detailed methodology for this combination, specifically for the parameters of a BN. In this paper, we illustrate the techniques using a case study of an ecological risk assessment problem.

1 INTRODUCTION

Bayesian networks (BNs) are graphical models composed of both qualitative and quantitative components. The qualitative component is the graphical model, which represents the causal relationships between key factors and final outcomes. Associated with the structure is a set of conditional probability distributions, which form the quantitative component of a BN. These describe the strength of relationships between linkages in the model, which can be estimated using a combination of sources.

BNs are increasing being used in ecological applications as they offer a pragmatic and scientific approach to modelling complex ecological systems where high uncertainties (aleatory and epistemic) exist. Unlike many other ecological modelling approaches, BNs can utilise prior knowledge and data to model systems. Furthermore, BN models are particularly useful for analyzing and communicating causal assumptions not easily expressed using mathematical notation, and for analyzing multivariate and complex relationships among variables.

Despite BNs having the ability to combine information from multiple sources, many reported BN applications to date, including non-ecological applications (see Korb and Nicholson (2004) for a recent survey), have been built through knowledge elicitation from domain experts (DEs) only. In general, this is both a difficult and time consuming task, with problems involving incomplete knowledge of the domain, and common human difficulties in specifying and combining probabilities. Therefore, recent interest has focused not only on combining parameter estimates from both DEs and data, but also developing a process where the efforts of the DE are focused on the more influential or 'sensitive' parts of the BN model.

Thus far, a methodology and associated support tools for Knowledge Engineering Bayesian Networks (KEBN) are not well developed. Spiral, prototype-based approaches to KEBN have been proposed (e.g., Laskey and Mahoney (2000); Korb and Nicholson (2004)), based on successful software development processes. However, these provide little guidance on integrating the knowledge engineering of the qualitative and quantitative components or again on how to combine knowledge elicitation from DEs and automated knowledge discovery methods. While there have been attempts at the latter, they remain rudimentary (e.g., Onisko et al. (2000); Nicholson et al. (2001)).

In Woodberry et al. (2004a), we presented a more detailed methodology, based on the spiral prototype model, for knowledge engineering the quantitative component of a BN. Our methodology explicitly integrates KE processes using both DEs and machine learning, in both the parameter estimation and the evaluation phases. The methodology was developed during the knowledge engineering of an ecological risk assessment (ERA) application (Section 3). This paper illustrates the methodology using this case study.

The focus of the ERA application was the development of a decision support tool to aid in fisheries management in the Goulburn Catchment, Victoria, Australia. After specification of the BN model structure, only limited datasets and knowledge elicitation information from DEs was available for model parameterization. The parameterisation of the ERA application using our spiral knowledge engineering methodology (Woodberry et al., 2004a) is described in this paper.

2 BAYESIAN NETWORKS

2.1 Background

A BN is a graphical representation of a joint probability distribution over a set of statistical variables. The structure is a directed acyclic graph (DAG), which is made up of a collection of nodes that represent variables.

Associated with variables are ranges of states and a conditional probability table (CPT), which describes the probability of each value of the child, conditioned on every possible combination of values of its parents. If a variable has no parents, it is described by a marginal probability distribution. Given both the qualitative and the quantitative parts, probabilities of any query variables posterior to any evidence can be calculated.

A prior (unconditional) probability represents the likelihood that a variable will be in a particular state and the conditional probability calculates the likelihood of the state of a variable given the states of input variables. The posterior probability distribution for a variable is calculated when given a new set of observations. Thus, BNs exploit the distributional simplifications of a network structure by calculating how probable events are, and how these probabilities can change given subsequent observations or predict change given external interventions (Borsuk and Reckhow, 2004).

2.2 Quantitative Knowledge Engineering Methodology

A methodology, presented in Woodberry et al. (2004a), for the quantitative knowledge engineering of BNs is outlined in Figure 1. This method illustrates flows (indicated by arrows) through the different KE processes (rectangular boxes), which will be executed either by humans (the DE and the knowledge engineer, represented by clear boxes) or computer programs (shaded boxes). Major choice points are indicated by hexagons.

The initial stage in the development spiral is Structural Development and Evaluation, which on the first iteration will produce an unparameterized causal network.

The next step in Figure 1 is parameter estimation, which involves specifying the CPTs for each node. The parameter estimates can be elicited from DEs, which can also include the domain literature as a source of parameter estimates. Parameter estimates can also be learned from data (path 2) or, as proposed here, generated from a combination of both sources (an example is shown in path 3). As BN development is an iterative process, in early prototypes the parameter estimates need not be exact, and uniform distributions can be used if neither domain knowledge nor data are readily available.

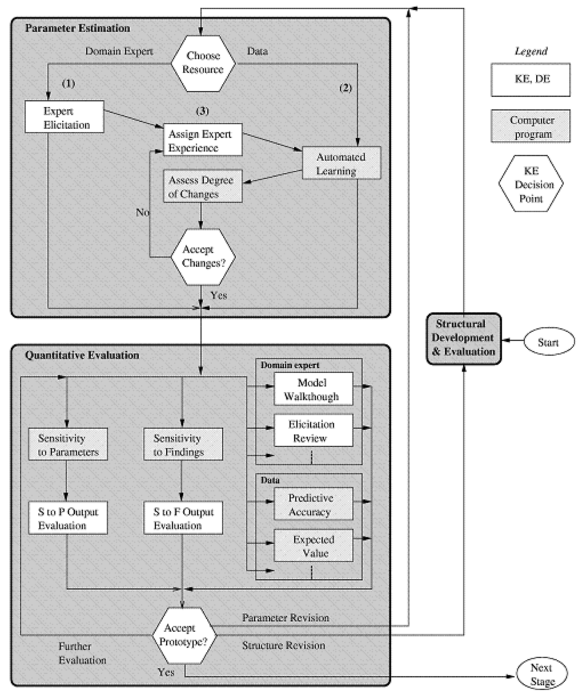

Figure 1: Knowledge Engineering Methodology

The second major aspect in the KEBN process is quantitative evaluation. Evaluative feedback can be generated using either DEs or data or both, as we have done here. When data is available, several measures can be used to evaluate BNs, including predictive accuracy, expected value computations and information reward. DE evaluation techniques include elicitation reviews and model walkthroughs (see Figure 1). Another kind of evaluation is sensitivity analysis.

Sensitivity analysis involves analyzing how sensitive the network is, thus determining how responsive probabilities of query nodes are to changes in parameters and inputs. Measures for these can be computed automatically using BN tools (shown as Sensitivity to Parameters and Sensitivity to Findings processes, in Figure 1), but these need to be evaluated by the DE in conjunction with the KE.

A more comprehensive description of the KE process is provided in (Woodberry et al., 2004a).

3 CASE STUDY: ERA APPLICATION

In the ecological risk assessment domain, there is an increasing need for environmental decision support tools that are able to model complex ecosystems in an integrated framework, acknowledging that uncertainties in input information and predictive outputs exist.

Currently, few tools meet these requirements. Those that are available are often highly complex, poorly tractable and not particularly user friendly.

The objective of the Goulburn Catchment (Victoria, Australia) ecological risk assessment described in this paper was to support future decision-making in the catchment by developing a predictive model with the following requirements:

- Quantify linkages between system variables and native fish communities in a highly complex system;
- Incorporate knowledge elicited from DEs;
- Incorporate existing datasets;
- Identify key risks to native fish communities;
- Predict changes to fish communities given system changes;
- Communicate uncertainty in predictions.

The BN modelling approach was trialed to determine if it could meet these requirements.

3.1 Background: Goulburn Catchment

The main stem of the Goulburn Catchment, the Goulburn River, is the largest tributary of the Murray-Darling Basin in the State of Victoria (Australia). The lowland Goulburn River extends from Eildon to its confluence with the Murray River at Echuca (Figure 2). Many rivers and creeks enter the 436 km lowland stretch of the Goulburn River.

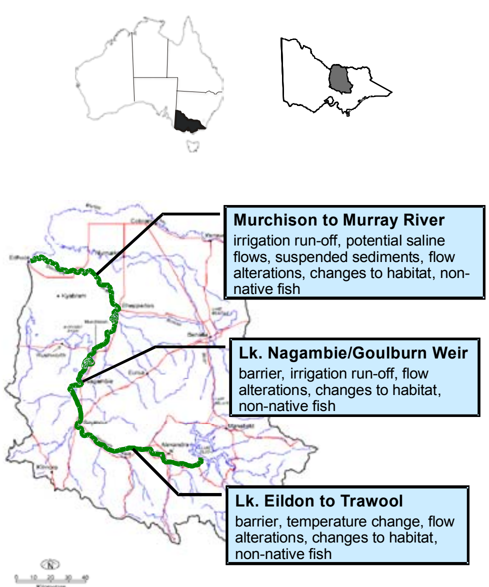

Figure 2: Location of Goulburn Catchment, showing major system changes

The headwaters of the Goulburn River flow into Lake Eildon. Water released from Lake Eildon is delivered 218 km downstream to Goulburn Weir. From Goulburn Weir, outflows are to the lower Goulburn River and three irrigation channels.

There is evidence that native fish communities in the Goulburn Catchment have declined over the past 100 years (see Pollino et al. (accepted)). Four major factors have been identified as influencing native fish abundance and diversity in the Goulburn River, being water quality, flow alterations, in-stream habitat and biological interactions. Although the processes and interactions between these factors and their link to native fish decline are broadly understood, quantitative models to assist in environmental management do not exist.

3.1.1 Graphical Structure

One version of the BN for this application is shown in Figure 3. For the purposes of this paper, variable names in Figure 3 have been simplified.

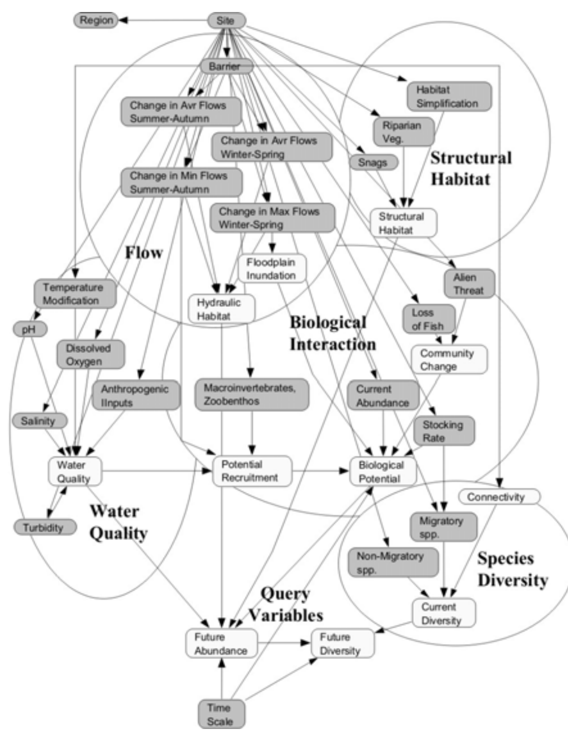

Figure 3: ERA BN Structure

The structure of the BN is based on a comprehensive conceptual model developed by DEs. This model consists of five interacting components - water quality, hydraulic habitat, structural habitat, biological potential and species diversity. The model has two query variables: Future Abundance and Future Diversity.

The query variables were also established in collaboration with DEs and are based on a very pragmatic management need in the Goulburn Catchment, being to assess what conditions are required to establish sustainable native fish communities. Given that no native fish recruitment data was available, the surrogate endpoints used were based on what information was available.

To assess the impacts of human-related activities on native fish communities, it was important to begin constructing the BN by establishing linkages between processes and activities of importance. Clearly, the Goulburn Catchment is a highly complex system with multiple factors interacting and influencing fish communities. The empirical relationships within and between chemical, physical and biological system components have not been previously characterized.

A temporal scale of one and five years are considered in the model. The spatial scale considers 23 sites in the catchment, which are further aggregated into 6 regions.

The development of the model structure is not described in any more detail in this paper, other than to acknowledge that it was undertaken via an iterative process.

4 PARAMETER ESTIMATION

To parameterise the ERA BN, all variables were discretised into sub-ranges. Although discrete variables have a finite set of possible states, continuous data can be entered into these variables. Where possible, variables were discretised using classifications and thresholds from existing management guidelines (e.g. Water Quality targets).

After discretisation, variables were split into two groups: those with data were initially given uniform probability distributions; the remainder were elicited from DEs (including DE literature). In Figure 3, data input variables are indicated by shading, although some sites have data missing for particular variables. Missing data was predominately due to an absence of monitoring data. Where variables had only limited or no data available, parameters were initially elicited from DEs (see Section 4.1). These variables are indicated by lighter shading.

4.1 Elicitation from Experts

The DEs who participated in this study included ecology experts and natural resource managers. Elicitation was conducted in a workshop environment and individually.

Direct elicitation of DEs employs such questions as ``What is the probability that variable A takes this state given these parent values?" This can employ the use of use frequencies, odds, or qualitative elicitation, using terms such as 'high' or 'unlikely', with the mapping to actual probabilities calibrated separately. When eliciting precise parameters, it can also be useful to elicit an acceptable range for the parameter. Intervals can be used during later evaluation to identify parameters needing further attention, as we shall see.

DEs were asked to report their confidence in these estimates, which was categorized as either low or high. DEs tended to be more confident estimating variables pertaining to the physical and chemical relationships in the system and less so with the biological relationships. In this study, the elicited confidence applied to the node, i.e., to the whole CPT, rather than to individual parameters; although, this need not be the case in general.

4.2 Learning Parameters from Data

Data is available for all 23 sites in the BN, which were arranged into a series cases (based on the date of a fish survey) for data learning. The number of cases per site ranged from between 3 and 272. In total, there were 949 cases. For the purpose of parameter learning each case was counted twice to match cases with one- and five-year projections of future abundance and future diversity.

In circumstances where data is of good quality and voluminous for parameter estimation (e.g. data only nodes in Figure 3), or where DE estimates are to be supplemented with data, 3 automated algorithms are available to do this in the Netica BN software (Norsys, 2000). These include: the Lauritzen Spiegelhalter method (LS) (Lauritzen and Spiegelhalter, 1990); the expectation maximization (EM) algorithm (Dempster et al., 1977); and the gradient descent (GD) algorithm (Norsys, 2000). Missing data is dealt with by finding the parameterizations which yield the greatest likelihoods given the data available.

Of the LS, EM and GD automated learning methods available, the EM method was selected for the ERA BN, since the LS method was not very useful with many parent instantiations missing in the data and the GD method was susceptible to local maxima. Automated learning trials were then carried out using EM in order to investigate the effects of different weightings of expert elicited CPTs. A pre-trial with the LS method was used for comparative purposes.

4.3 Combining DE and Quantitative Estimations

When combining elicitation and data-based parameterizations, elicited information is weighted relative to the data available. In Figure 1 this is done in the Assign Expert Experience process, where an experience weighting is assigned to the expert parameter estimates, based on the confidence in the estimates obtained during expert elicitation. These are then treated as equivalent to the size of a hypothetical initial data sample (the equivalent sample size, or ESS).

After incorporating the data in parameter estimation, the next step is to compare the new CPT with the original. In Figure 1 we consider this to be an automated process, Assess Degree of Changes. Parameters estimated from the data that are outside acceptable range of values elicited by experts be flagged for attention. An alternative method for comparing the parameterizations is Bhattacharyya distance

(Battacharyya, 1943), which measures the distance between the two probability distributions. This distance is computed for each possible combination of parent values; higher distances between conditional distributions trigger further attention. The DE must then assess whether these flagged parameter refinements obtained after automated learning are acceptable (in the Accept Changes decision point in Figure 1).

Table 1. Nodes whose CPTs were first expert elicited, with the different experience weightings used for trials of the EM automated learning method.

	H=10,M=5 Trial No.			H=20,M=10 Trial No.			Combined Trial No.	
Node	1	2	3	1	2	3	4	5
Water Quality	10	15	18	20	25	25	25	24
Hydraulic Habitat	10	15	18	20	25	25	25	24
Structural Habitat	10	7	4	20	15	10	1	1
Biological Potential	5	2	4	10	5	5	5	5
Temperature Modification	10	5	1	20	15	10	1	1
Community Change	5	1	1	10	5	1	1	1
Floodplain Inundation	10	5	1	2-	15	10	1	1
Potential Recruitment	5	1	3	10	5	2	3	3
Connectivity	10	10	10	20	17	14	12	12
Migratory spp	5	10	15	10	15	15	15	16
Current Diversity	5	1	1	10	7	4	1	2
Future Abundance	5	2	4	10	7	4	5	6
Future Diversity	5	5	5	10	5	5	5	5
Remaining Nodes	0	0	0	0	0	0	0	0

When using the EM method to refine the parameters of all nodes in the ERA BN, a series of trials were conducted (see Table 1). Each trial used a series of experience weightings. Each of these trial EM parameterizations was compared, using the Bhattacharya distance, with the LS BN, and an assessment was made as to whether the degree of change was acceptable. If the change was deemed unacceptably large, the ESS was increased, while if there was no or minor changes, the ESS was decreased. This assessment process was iterated, comparing the new EM parameterization with the LS parameterization and setting a new ESS value, $W_{i+1,}$ using the Algorithm shown below.

ALGORITHM: Adjusting ESS

Loop until ESS values converge
Parameterize network with current ESS values
Switch
Case changes unrealistic: $W_{i+1} \leftarrow W_i$ *+ upLarge (5)*
Case would allow greater changes : $W_{i+1} \leftarrow W_i$*+ upSmall (3)*
Case little OR no change: $W_{i+1} \leftarrow W_i$ *+ downLarge (5)*
Case insignificant change: $W_{i+1} \leftarrow W_i$ *- downSmall (3)*
Case changes become unrealistic: $W_{i+1} \leftarrow W_{i-1}$ *+ bounceup (2)*
Case changes disappear: $W_{i+1} \leftarrow W_{i-1}$ *- bouncedown (2)*
Case final trials, small adjustments needed: $W_{i+1} \leftarrow W_{i-1} \pm$ *tweak (1)*
Case changes acceptable: $W_i \leftarrow W_i$
End Loop

5 QUANTITATIVE EVALUATION

After parameterization of the BN, the second major aspect of quantitative knowledge engineering is evaluation (Figure 1), which guides further iterations of BN development.

5.1 Evaluation using Data

When data is available, it can be used for evaluation. Where the data is also being used to learn the structure or the CPTs, it is necessary to divide it into training data and test data, so that evaluation is not done with the very same data used for learning. The most common method of evaluation is to determine the predictive accuracy of the BN, which measures the frequency with which the modal node state (that with the highest probability) is observed to be the actual value.

In the ERA case study, case data was randomly split so that 80% of data were used for training and 20% used for testing.

The error rate of the query nodes, Future Abundance and Future Diversity, were only 5.8% and 0%, respectively. The low error rates reflect the lack of variability available in the dataset, and thus the full spectrum of variability in the network cannot be tested. Therefore, as data was limited, automated evaluation was of limited use. Other, less formal, data evaluation was conducted and is discussed in Woodberry et al. (2004b).

5.2 Evaluation using DE

Even when adequate data is available, it is important to involve the DE in evaluation. If expert elicitation has been performed, a structured review of the probability elicitation is important. This procedure could involve: comparing elicited values with available statistics; comparing values across different DEs and seeking explanation for discrepancies; double-checking cases where probabilities are extreme (i.e., at or close to 0 or 1), or where the DEs have indicated a low confidence in the probabilities when originally elicited.

For the ERA BN, the domain expert developer conducted a semi-formal model walkthrough with ecology experts and natural resource managers, with positive feedback. It is recognized that to strengthen the model more case data is needed.

5.3 Sensitivity Analysis

Sensitivity analysis is used to measure how sensitive the network is, in terms of changes in updated probabilities of some query nodes to changes in parameters and inputs. We review two types of sensitivity analysis. One type, "sensitivity to findings," looks at how the BN's posterior distribution changes under different observed conditions. The other, "sensitivity to parameters," looks at how the model's distribution changes when particular parameters are altered. Curiously, researchers thus far appear to have employed one or the other of these methodologies, but not both in any one study (e.g., Laskey and Mahoney (2000); Rieman et al. (2001); Coupe and Van der Gaag (2002)). Both are needed for a careful and thorough investigation of the properties of a network.

In the ERA BN, sensitivity analyses were used to identify variables that were highly sensitive to change, so as quantification efforts in the proceeding model iterations were focused. Where the accuracy of parameters could not be improved, they were investigated to determine if they represented knowledge gaps, indicating where further monitoring and research efforts are required.

Variables that were identified as contributing little to improving the predictive accuracy of the model were given less attention.

5.3.1 Sensitivity to Findings

In BNs, the properties of d-separation can be used to determine whether evidence about one variable may influence belief in a query variable. It is possible to measure this influence and rank evidence nodes by how much of an effect they have. This information can be used to provide guidance for collecting the most informative evidence or as a check on whether the model reflects the DE's intuitions.

Sensitivity to findings can be quantified using two types of measures, entropy and mutual information. Entropy, H(X), is commonly used to evaluate the uncertainty, or randomness, of a probability distribution:

$$H(X) = - \sum_{x \notin X} P(x) \log P(x) \tag{1}$$

Measuring the effect of one variable on another is referred to as the mutual information (MI):

$$I(X|Y) = H(X) - H(X|Y) \tag{2}$$

We have implemented this type of sensitivity to findings (see Woodberry et al. (2004b)). Our algorithm computes and displays both the entropy of a specified query node and the ranked mutual information values for a specified set of interest nodes, given a set of evidence for some other observed nodes. The user can subsequently investigate how changes to the evidence will affect the entropy and MI measures. This process allows the DE to identify whether a variable is either too sensitive or insensitive to other variables in particular contexts, which in turn may help identify errors in either the network structure or the CPTs.

5.3.2 Sensitivity to Parameters

Sensitivity analysis could be performed using an empirical approach, by altering each of the parameters of the query node and observing the related changes in the posterior probabilities of the target node. However, such a straightforward analysis can be extremely time consuming, especially on large networks. Coupe and Van der Gaag (2002) address this difficulty by first identifying a "sensitivity set" of variables given some evidence.

A sensitive set of variables are those that can potentially change. Thus, less effort can be spent on the remaining variables. The sensitivity set can be found using an adapted d-separation algorithm (see Woodberry et al. (2004b)). Coupe and Van der Gaag (2002) also demonstrated that the posterior probability of a state, given evidence under systematic changes to a parameter value, can be given a functional representation, either linear or hyperbolic.

For the ERA BN, we implemented sensitivity to parameters (see Woodberry et al. (2004b)). When a particular evidence instantiation is set, our algorithm (see Woodberry et al. (2004b)) identifies the type of sensitivity function for the parameters by checking whether the query node has any observed descendant nodes. Once the sensitivity function is determined for a parameter, its coefficients can be computed. If the plotted sensitivity function does not behave as the DE expects (its slope, direction or range is unexpected), then this could indicate errors in the network structure or CPTs.

The revised normalized probability distribution of the test node is set by first selecting a new value, P_{new} for the parameter under investigation, Pj. The remaining parameters, Pi, are normalized to retain relative values by the updating function,

$$Pi \Leftarrow Pi \times \frac{1 - Pnew}{1 - Pj}, i \neq j \tag{3}$$

before the parameter under study is updated,

$$Pj \neq Pnew \tag{4}$$

In the ERA BN, assessments for the network's conditional probabilities were systematically varying over a plausible interval, being 0 to 1 used in this study. The effects on the behaviour of the system were examined.

Results of sensitivity to parameters indicate that the network parameters are relatively insensitive to change. The most sensitive scenario was identified as: probability that future abundance is low given, water quality was low, structural habitat was low, biological potential was low, and hydraulic habitat was low at Eildon over a one year time scale (Figure 4). These environmental conditions represent the "worst case scenario" for native fish in the catchment.

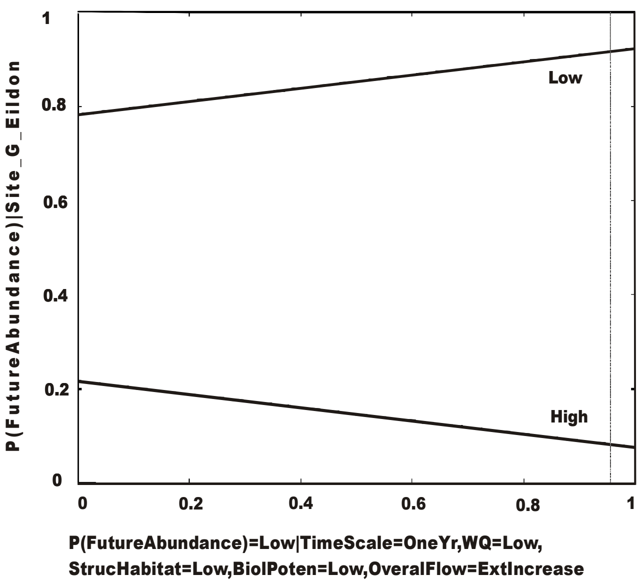

Figure 4: Sensitivity to parameters output showing slope of change for high and low future native fish abundances at one site.

5.4 Model predictions

Fisheries data from each site in the catchment were plotted against model predictions (Figure 5). The model predictions are based on existing environmental conditions in the Goulburn Catchment.

Comparisons between data and predictions are only relative due to the different scales (probability versus relative abundances); however, trends between the fisheries dataset and the predictive outputs are generally maintained.

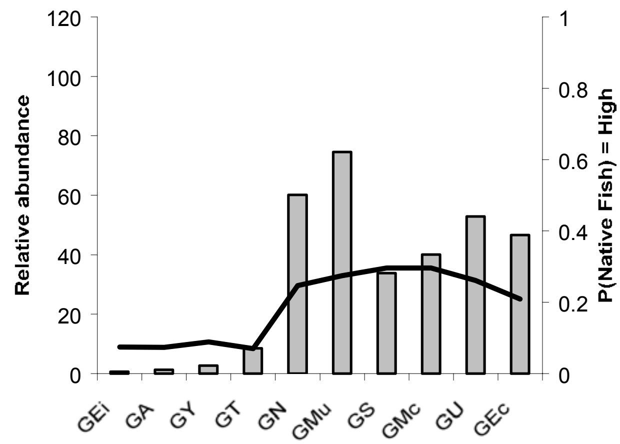

Figure 5: Relative abundance data (left axis - bars) versus BN model predictions (right axis - line) for sites in the Goulburn main channel.

Clearly, fish communities in the upper part of the catchment (GEi, GA, GY, GT) are under stress. Using the results of the sensitivity analyses (not shown), at these four sites, water quality (temperature) and changed hydrology appear to be the variables primarily influencing native fish abundance. At the remaining sites, fish abundance is primarily under the influence of the biological potential (potential recruitment and current abundance), water quality (turbidity, dissolved oxygen and pH), and flow.

6 ACCEPT MODEL PROTOTYPE

Quantitative evaluation can be used to identify problems with the BN structure and parameters. After the model has been evaluated using a particular technique, the KE and DE must determine whether the prototype is to be accepted for the next stage of development. This decision is not intended to be the end of the knowledge engineering, or even prototyping, process.

If the prototype is not sufficiently validated for prototype acceptance, Further evaluation is one option for the KE and DE. It will often be necessary to use multiple evaluation techniques to validate the model: for example, sensitivity to findings and parameter analyses evaluate different aspects of the model with little overlap, and hence don't substitute for each other. If problems with either the structure or the parameters have been identified, it will be necessary to re-visit the relevant KE processes, Structural Development Evaluation or Parameter Estimation respectively, via the main spiral iteration in Figure 1.

7 DISCUSSION

7.1 Knowledge Engineering

In this paper we have described the use of a methodology for combining expert elicitation and data for parameterision of BNs, an important research topic that has been widely acknowledged in the BN field but little developed.

In many ecological applications, including our ERA case study, information sources can be considered to be poorly documented, poorly understood, and generally incomplete. Although other causal network structures (e.g. Borsuk and Reckhow (2004)) have been developed using such information sources, unlike this study, parameter estimates for a variable were obtained from only one source (i.e. DEs or data). To parameterise the ERA BN model, we directed our efforts towards combining multiple information sources, each with associated uncertainties, and undertaking an iterative process to derive acceptable parameter estimates.

Evaluative methods, including sensitivity analyses, were used to investigate the uncertainties and inaccuracies in model structure, relationships and outputs (Coupe and Van der Gaag, 2002). This process enabled a more targeted approach in identifying parameters that needed to be accurately quantified. Thus, based on these results, recommendations for targeted monitoring and studies can be made.

By utilizing a BN knowledge engineering spiral, we developed a model prototype that has been accepted for use in an ecological risk assessment domain. Future studies will investigate further testing and refining this methodology in other domains, and to iteratively assess and develop the deployed ERA model.

7.2 ERA Application

The development of quantitative decision-support systems in environmental management are considered to be of high priority today. By using predictive models, it is anticipated that decisions will be more robust, defensible and tractable. However, given that both the understanding of many complex ecological systems is considered to be limited, and the existing modelling technologies for describing such systems are poor, progress has been limited.

In this study, the BN modelling technology was trialed in an effort to model a complex ecological system, the Goulburn catchment. As with many catchments, there were high uncertainties associated due to the lack of knowledge of the relationships between variables and lack of data available. Nonetheless, the ERA BN developed has the ability to predict the abundance and diversity of native fish communities based on existing and predicted changes to environmental conditions. The model has the ability to assist in determining what management options are most favorable for maintaining and rehabilitating fish communities at multiple spatial scales. However, the model requires further testing in the field to determine its accuracy pre- and post- management interventions or system changes.

As new information is made available, this can be incorporated into the model. Unlike many other modelling approaches, BN can continually be developed, adapted and refined. This process can be conducted using the parameterization and evaluation process described in this paper. Consequently, BN models have the potential to become both an important and adaptive learning tool, as well as an important adaptive resource management tool.

Further investigations are underway to make the ERA BN specific to different types of fish communities and to trial it in different catchment areas.

REFERENCES

Battacharyya, A., 1943. On a measure of divergence between two statistical populations defined by their probability distributions. *Bull. Calcutta Math Soc.,* 35: 99-110.

Borsuk, M.E., Stow, C.A. and Reckhow, K., 2004. A Bayesian network of eutrophication models for synthesis, prediction and uncertainty analysis. *Ecol. Modelling,* 173: 219-239.

Coupe, V.M.H. and van der Gaag, L.C., 2002. Properties of sensitivity analysis of Bayesian belief networks. *Ann. Math. Art. Intell.,* 36: 323-356.

Dempster, A., Laird, N. and Rubin, D., 1977. Maximum likelihood from incomplete data via the EM algorithm. *J. Royal Stat. Soc. B*, 39: 1 - 38.

Korb, K.B. and Nicholson, A.E., 2004. Bayesian Artificial Intelligence. Chapman and Hall / CRC Press, London, 364 pp.

Laskey, K.B. and Mahoney, S.M., 2000. Network Engineering for agile belief network models. *IEEE Trans. Know. Data Eng.,* 12: 487-498.

Lauritzen, S.L. and Spiegelhalter, D.J., 1990. Local computations with probabilities on graphical structures and their application to expert systems. In: G. Shafer and J. Pearl (Editors), Readings in Uncertain Reasoning. Morgan Kaufmann, pp. 415-458.

Nicholson, A.E., Boneh, T., Wilkin, T., Stacey, K., Sonenberg, L. and Steinle, 2001. A case study in knowledge discovery and elicitation in an intelligent tutoring application. In: Breese and Koller (Editors), UAI01, Seattle, pp. 386-394.

Norsys, 2000. Netica. www.norsys.com.

Onisko, A., Druzdel, M.J. and Wasyluk, H., 2000. Learning {B}ayesian network parameters from small data sets., Workshop on Bayesian and Causal networks, ECAI-2000.

Pollino, C.A., Feehan, P., Grace, M. and Hart, B., accepted. Fish communities and habitat changes in the highly modified Goulburn Catchment, Victoria, Australia. *Marine and Freshwater Research.*

Rieman, B.E., Peterson, J.T., Clayton, J., Howell, P., Thurow, R., Thompson, W. and Lee, D.C., 2001. Evaluation of potential effects of federal land management alternatives on trends of salmonids and their habitats in the interior Columbia River basin. *Forest Ecology and Management.,* 153: 43-62.

Woodberry, O., Nicholson, A.E., Korb, K.B. and Pollino, C.A., 2004a. Parameterising Bayesian Networks, To appear in *Proceedings of the 17th Australian Joint Conference on Artificial Intelligence*, December 2004, Cairns.

Woodberry, O., Nicholson, A.E., Korb, K.B. and Pollino, C.A., 2004b. A methodology for parameterising Bayesian networks, Technical Report, School of Computer Science and Software Engineering, Monash University.

BIOGRAPHIES

CARMEL POLLINO Ph.D. earned her doctorate in ecotoxicology at RMIT University (2000). After working on conducting laboratory investigations in ecological risk assessment at the City University of Hong Kong, Carmel joined the Water Studies Centre at Monash University as a Research Fellow. Her research focus is on the use of Bayesian networks for risk assessments and the wider adoption of risk assessment into environmental management.

OWEN WOODBERRY received his Bachelor of Computer Science with Honours at Monash University in 2003. He has also worked as a Research Assistant with the Water Studies Center continuing the work of his Honours project. He is currently undertaking his PhD in Evolutionary Ethics with the Computer Science Department also at Monash University. His research interests include Bayesian Networks, Ecology, Ethics, Morality, Sociobiology, ALife and Artificial Intelligence.

ANN NICHOLSON D.Phil., did her undergraduate computer science studies at the University of Melbourne and her doctorate in the robotics research group at Oxford University (1992), as a Rhodes Scholar, working on dynamic Bayesian networks for discrete monitoring. She then spent two years at Brown University as a post-doctoral research fellow before taking up a lecturing position at Monash University in Computer Science. Her general research focus is AI methods for reasoning under uncertainty, while her current research includes knowledge engineering with Bayesian networks, applications of Bayesian networks and user modeling.

KEVIN KORB Ph.D., earned his doctorate in the philosophy of science at Indiana University (1992) working on the philosophical foundations for the automation of Bayesian reasoning. Since then he has lectured at Monash University in Computer Science, combining his interests in philosophy of science and artificial intelligence in work on understanding and automating inductive inference, the use of MML in learning causal theories, artificial evolution of cognitive and social behavior and modeling Bayesian and human reasoning in the automation of argumentation.