Issues with modelling decisions in Bayesian networks

David Albrecht Ann Nicholson

Monash University

Keywords: Decision Networks; Constraints; Shielding

ABSTRACT

Designing a Bayesian Network or a Decision network (influence diagram) for a particular domain is a difficult task. It involves many decisions about how to represent the domain knowledge, including which variables should be represented, what distributions the variables should have, and the connections between the variables. When the domain involves decision making, there are modelling choices to be made about how to represent the decisions and the factors influencing those decisions.
In this paper we will consider a simple example that involves these modelling considerations, that can be either represented as a Bayesian network (with or without explicit representation of decisions) or as a Decision network (with explicit decision and utility nodes). Through this example, we consider the effect of various modelling choices on belief updating and the faithfulness of the resultant network to the original problem.

1. INTRODUCTION

When modelling a domain that involves decision making one should consider using a Decision Network (influence diagram). However, before deciding whether to use a Decision Network it is important to understand the effect of various modelling choices on belief updating.

In this paper we will present a simplified version of an example presented by Crowley (2005). The problem involves the allocation of lecturers to teach a course with the constraint that at least one lecturer has to teach it. As we describe in Section 2, Crowley incorporated an explicit constraint node, which lead to unwanted side-effects; while he proposed a novel "shielding" method, this did not resolve the problem satisfactorily.

We will look at how the problem can be either represented as a Bayesian network (with or without explicit representation of decisions) or as a Decision network (with explicit decision and utility nodes). Through this example, we consider the effect of various modelling choices on belief updating and the faithfulness of the resultant network to the original problem. The case study also highlights some more general principles for Bayesian network modelling.

2. CROWLEY'S EXAMPLE

Crowley (2005) introduces an example problem that involves instructors (lecturers) being assigned to teach a course. We will consider a simplified version of this problem involving only two lecturers.

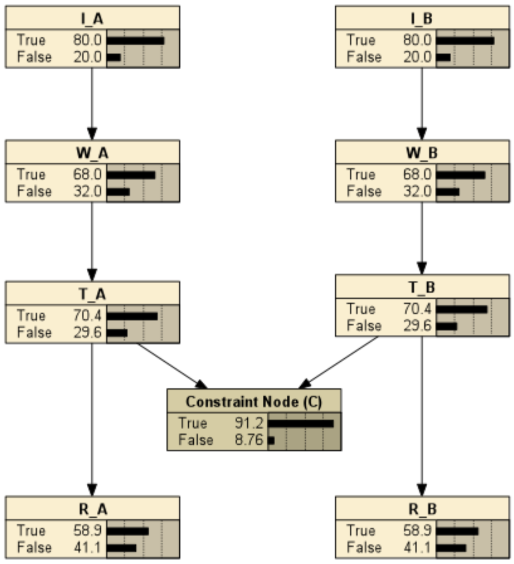

Figure 1: Lecturer example with constraint node not set.

In our version each lecturer is independently interested in teaching the course, which we represent by the variables $\mathbf{I_A}$ and $\mathbf{I_B}$ respectively. Whether a lecturer wants to teach the course, represented by the variables $\mathbf{W_A}$ and $\mathbf{W_B}$ respectively, depends only on their interest. Whether they teach the course, represented by the variables $\mathbf{T_A}$ and $\mathbf{T_B}$, only depends on whether they want to give the course.

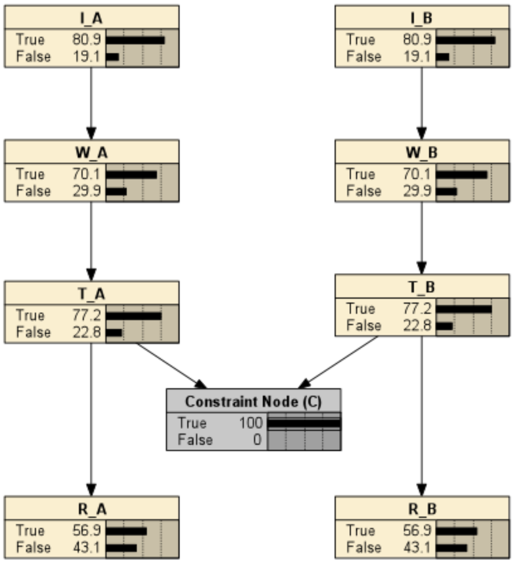

Figure 2: Lecturer example with constraint node set.

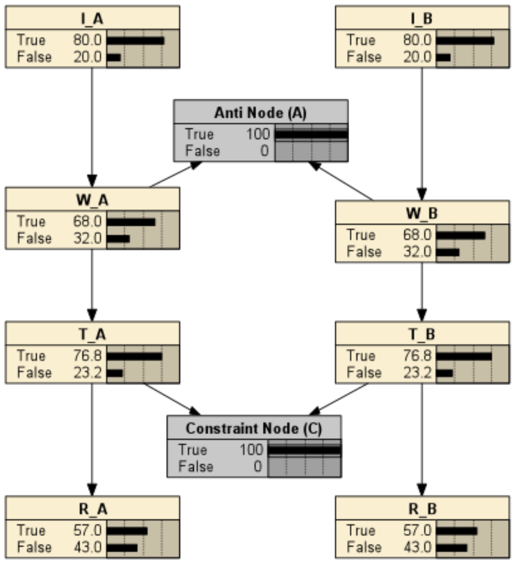

Figure 3: Lecturer example with anti node

Also, whether they reach their research goals, represented by the variables $\mathbf{R_A}$ and $\mathbf{R_B}$, depends only on whether they teach the course. Finally, we will assume that at least one of the lectures must teach the course.

Crowley modelled the constraint that at least one lecturer must teach the course in a variety of ways. One way (see Figure 1) involved introducing a so-called binary valued constraint node (Jensen and Nielsen, 2007, p.74), $\mathbf{C}$, whose parents are $\mathbf{T_A}$ and $\mathbf{T_B}$, and the value of this node is true if and only if one of the lectures does teach the course.

Table 1: Conditional Probability Table for **C**

$\mathbf{T_A}$	$\mathbf{T_B}$	$\mathbf{C}$
true	true	**true**
true	false	**true**
false	true	**true**
false	false	**false**

To model the constraint the value of **C** is then set to **true** (see Figure 2. Crowley noticed that this had the 'side-effect' that the interest in the course by the lecturers changes. In Figure 1, $p(\mathbf{I_A} = \mathbf{true}) = 0.8$ while in Figure 2, $p(\mathbf{I_A} = \mathbf{true}) = 0.809$. Crowley considered this effect inappropriate and so developed methods for modifying the network so that this effect would not occur.

The simplest approach (see Figure 3) he considered involved adding another binary valued node, **A**, called an anti-node to the network and setting its value to true. The parents of **A** are $\mathbf{W_A}$ and $\mathbf{W_B}$ and conditional probability table (CPT) for **A** is designed so that the joint probability of $\mathbf{W_A}$ and $\mathbf{W_B}$ is independent of the value of **A** and the value of **C**.

To obtain the values for of the CPT for **A**, we let

$$p(\mathbf{A}|\mathbf{w_A}, \mathbf{w_B}) = \frac{\mathbf{K}}{p(\mathbf{C}|\mathbf{w_A}, \mathbf{w_B})},$$

for all values $\mathbf{w_A}$, $\mathbf{w_B}$ of $\mathbf{W_A}$ and $\mathbf{W_B}$, respectively, and choose **K** so that all values of the CPT of **A** are between 0 and 1. It then follows that **A** has the desired properties and the resulting network, when **A** is set to **true**, does not depend on what value you have chosen for **K**.

However, introducing **A** changes the values of $p(\mathbf{T_A})$ and $p(\mathbf{T_B})$. In Figure 3 (with **Anti-node** representing **A**), $p(\mathbf{T_A} = \mathbf{true}) = 0.768$ while in Figure 2, $p(\mathbf{T_A} = \mathbf{true}) = 0.772$. Moreover, it can be shown, either by Bayes' Theorem or constructing a simplified version of the network Figure 2 with only the nodes $\mathbf{T_A}$, $\mathbf{T_B}$ and $\mathbf{C}$, that the values for $p(\mathbf{T_A})$ and $p(\mathbf{T_B})$ are correct in Figure 2. Therefore by introducing a node **A** to solve one issue we have created another issue.

Other approaches Crowley (2005) considered involved combining new networks, which he called Anti-networks. In all cases the idea was to shield part of the network, in this case the nodes $\mathbf{W_A}$ and $\mathbf{W_B}$, together with their ancestors, from the influence of the constraint node, **C**. However, in all cases the introduction of these components causes other side-effects.

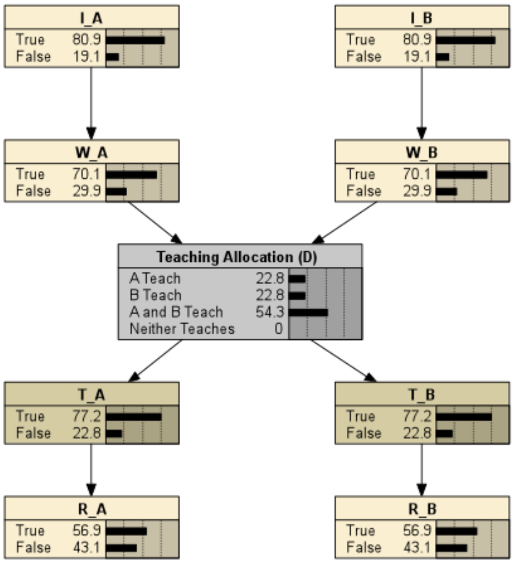

Figure 4: Lecturer example with the possibilities made explicit in a node

3. ALTERNATIVE BAYESIAN NETWORKS

We will now look at some alternative approaches to modelling the problem of the two lecturers given in the previous section.

Explicitly modelling all the alternatives

One approach involves introducing a node, **Teaching Allocation (D)**, which explicitly models the teaching allocation decision being made (see Figure 4). This node has parents $\mathbf{W_A}$ and $\mathbf{W_B}$, has children $\mathbf{T_A}$ and $\mathbf{T_B}$, and has the values **"A teaches"**, **"B Teaches"**, **"Both Teaches"**, **"Neither Teaches"**. To compute the CPT for **D** we use the the fact that:

$$p(\mathbf{t_A}, \mathbf{t_B}|\mathbf{w_A}, \mathbf{w_B}) = p(\mathbf{t_A}|\mathbf{w_A})p(\mathbf{t_B}|\mathbf{w_B}),$$

for values $\mathbf{t_A}, \mathbf{t_B}, \mathbf{w_A}, \mathbf{w_B}$ of $\mathbf{T_A}, \mathbf{T_B}, \mathbf{W_A}$, and $\mathbf{W_B}$, respectively. We remove the links between $\mathbf{W_A}$ and $\mathbf{T_A}$, and between $\mathbf{W_B}$ and $\mathbf{T_B}$. Finally, we add the negative evidence, that the value **"Neither Teaches"** is not possible.

This network is equivalent to the network with the constraint node (Figure 2). In fact it doesn't matter what way you model this constraint, if you do it properly you obtain an equivalent network. This can be seen by considering the joint probability of all the variables. So we still have the problem that the constraint alters the lecturers' interest in teaching the course.

More dependencies between variables

There are also other problems with the model. Consider the relationship in Crowley's original BN between the teaching of a course ($\mathbf{T_i}$) and whether they will achieve their research goals ($\mathbf{RA_i}$), represented by $\mathbf{T_i} \rightarrow \mathbf{RA_i}$. This structure implies that a lecturer's research output is the same whether teaching the course alone or jointly, which seems unlikely. It is straightforward to "fix" this problem by adding additional arcs from $\mathbf{T_A} \rightarrow \mathbf{RA_B}$ and $\mathbf{T_B} \rightarrow \mathbf{RA_A}$. However this does not resolve the more fundamental modelling problem, which is that there is a dependence between $\mathbf{T_A}$ and $\mathbf{T_B}$ which is not being captured in the network. If the variable of interest is not "Is lecture A/B teaching the course" but instead "Who is teaching the course", then we see that there should be a single variable **T**, with values "A teaches", "B teaches" and "Both A and B teach".[1] But if we collapse $\mathbf{T_A}$ and $\mathbf{T_B}$ into such a 3-valued node **T**, with $\mathbf{D} \rightarrow \mathbf{T}$, then we still have the problem of specifying a meaningless CPT row for the impossible situation $\mathbf{P(T|D = \text{"Neither teaches"})}$.

A simpler solution is to remove completely the **T** nodes, and have **Teaching Allocation (D)** be the single parent of both $\mathbf{RA_A}$ and $\mathbf{RA_B}$, as shown in Figure 5. Note that in this BN, we say that for each lecturer, the probability they will achieve their research goals given they are co-teaching the course is 0.65, half-way between the probability if they are teaching the course on their own (0.5) and not teaching the course (0.8).

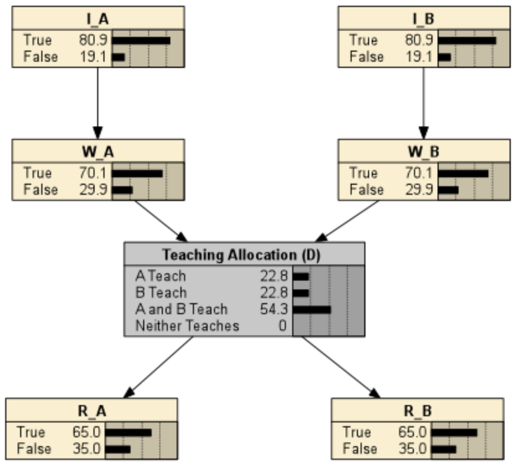

Figure 5: Lecturer example with the T nodes removed

[1] This is an example of Korb and Nicholson's (Korb and Nicholson, 2010) "Common Modelling Mistake 5: Separate nodes for different states of the same variable".

Explicitly modelling all influences

There is a further problem with both our alternative BNs containing the **Teaching Allocation** node: they show only that the teaching allocation depends on what the lecturers want, and do not capture explicitly the intuition that the decision will also take into account the likely impact on the lecturers research performance (an effect of that decision). Of course the $\mathbf{W_i}$ nodes, representing what the lecturers want, may be *implicitly* incorporating that preference, but that is captured only in the semantics of the node name, not in the network structure at all.

To solve this problem, and to ensure that the constraint in how the teaching is allocated does not change the lecturers' interest in teaching the course, you cannot model this with only a Bayesian Network. You need to include something else in your model.

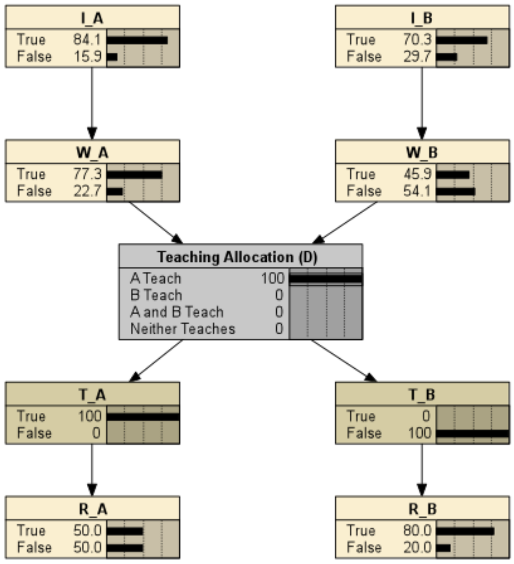

Figure 6: BN from Fig. 4 with observation that A has been allocated the course

4. MODELLING DECISION MAKING

Observations or interventions?

Crowley's original BN, and our variations, all provide a probabilistic model of the teaching allocation decision. When evidence is entered for those allocations, for example for $\mathbf{D}$ in our alternative BN, this changes both the predictions about whether the research s will be met ($\mathbf{R_A}$ and $\mathbf{R_B}$), and also change the beliefs in the lecturers interest (Figure 6). For the *observation* that someone has allocated the person a course (an observation) tells us that it is morely they wanted the course, which in turn means they were more likely to be interested in it. If the aim is to model decisions as causal *interventions*, then simply entering evidence into an ordinary BN is incorrect. This leads to the unwanted "side-effects" in ancestor nodes (as we have seen in Crowley's teaching example), which in turn can lead to incorrect predictions if there is an alternative path from the ancestor nodes to descendant nodes. See (Korb and Nicholson, 2010, Sec 3.8) for a more detailed presentation of this issue.

Alternative ways of modeling causal interventions in BNs have been proposed in the literature: Pearl's do-Calculus Pearl (2000) involves cutting the arcs to parent nodes to prevent the propagation back up the networks, while Korb et al. (2004) suggest an additional intervention node that allows modelling of probabilistic interventions and effectiveness (see also Korb (2011) in these proceedings). Moreover, there is an extension to Bayesian networks, so-called "decision networks" which not only model the probabilities associated with interventions, but support decision making by combining probabilities with the cost and benefits that influence decisions.

Decision networks

Decision networks (also known as "Influence Diagrams" Howard and Matheson (1981)) consist of 3 kinds of nodes: *chance nodes* (oval), which are the nodes in an ordinary BN, representing the domain variables; *decision nodes* (rectangular), representing the actions or decisions that may be taken; and *utility nodes* (diamonds) which model the utility associated with their parents nodes (chances nodes or decision nodes). The key point is that decision nodes model interventions in the system, and do not propagate changes in beliefs to their ancestors. Decision networks are used to compute the *expected utility* (**EU**) of each decision, combining the probability of each outcome, x, with the utility of that outcome and the decision:

$$EU(D = d) = \sum_x P(x|D = d)U(x, d)$$

A decision network for teaching allocation

Let us now recast Crowley's teaching allocation BN into a decision network for teaching allocation. There are only three possible decisions: **"A teaches"**, **"B teaches"** and **"Both A and B teach"** the course. These become the possible values for the decision node. This means there is no need to model explicit the "constraint" of the earlier BN models, that **"Neither A nor B teach"**, which is merely one of the many (possibly infinite!) alternatives that are not being considered.

The next step in the modelling process is to consider what the utilities are for this problem. There are clearly two main aspects: (1) it is "good" if lecturers are happy with their teaching allocation and (2) it is "good" if lecturers can meet their research goals.

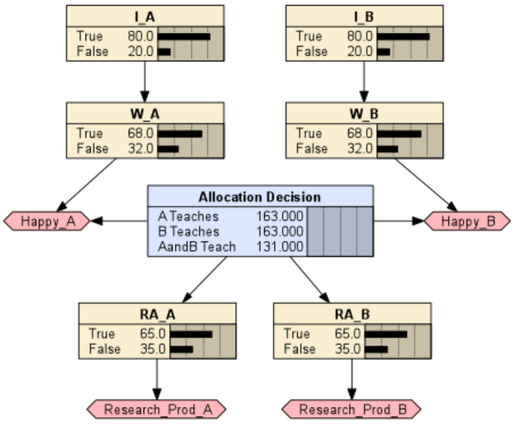

Figure 7: Decision network for the teaching allocation example showing posterior probabilities for the chance nodes, and the expected utilities for each decision.

Figure 7 shows a decision network for the teaching allocation example incorporating the decision **Allocation Decision** and four utility nodes: the "happiness" of the lecturers depending on whether they go the allocation they wanted[2] and the **Research_Prod** nodes. For illustration purposes, we have used the totally arbitrary utility function shown in Table 2.

Table 2: Utility tables for the decision network (the table for $\mathbf{Happy_B}$ is symmetric to $\mathbf{Happy_A}$)

$\mathbf{W_A}$	**Allocation Decision**	$\mathbf{Happy_A}$
true	A Teaches	**100**
true	A&B Teach	**50**
true	B Teaches	**0**
false	A Teaches	**-100**
false	A&B Teach	**-50**
false	B Teaches	**100**

$\mathbf{RA_i}$	$\mathbf{Research\_Prod_i}$
true	**100**
false	**-50**

Figure 7 also shows the posterior probabilities and the expected utilities for the decision network with no evidence added. We can see that, given the priors only, the "best" decision (that is, the one that maximises the expected utility) is to allocate both A and B to this course. Figure 8 shows three different scenarios for A and B's interest in the course: (a) both interested, (b) neither interested and (c) one interested and one not interested. We can see that the evidence effects the decision only through the impact on the utility, as the probabilities for the $\mathbf{T_i}$ and $\mathbf{RA_i}$ nodes are unchanged (as the decision hasn't been made yet). Conversely, Figure 9 shows what happens when the decision is taken: the impact is *only* on the $\mathbf{T_i}$ and $\mathbf{RA_i}$ nodes, with the **I** and **W** nodes unaffected.

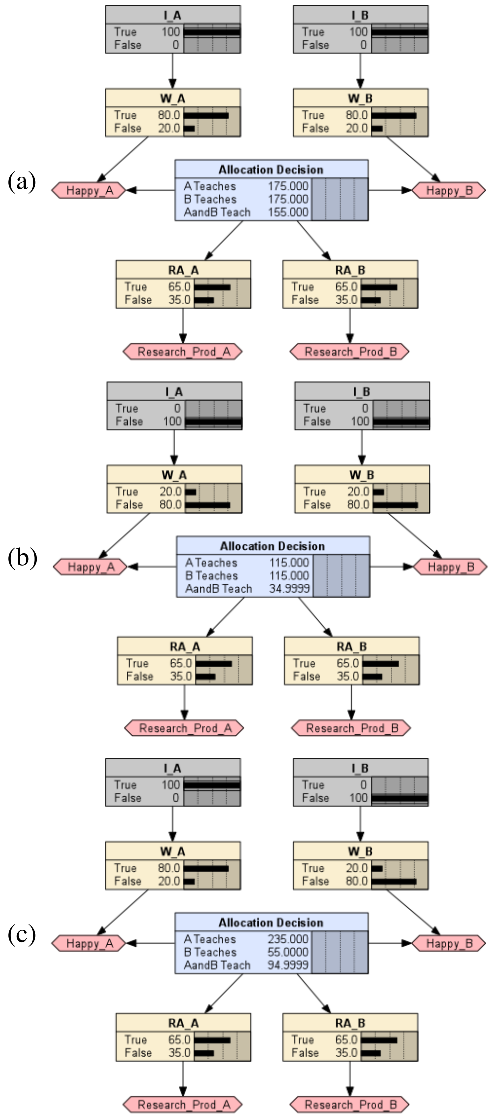

Figure 8: Three scenarios of lecturer interest, which impact on the expected utilities of the 3 alternative decisions.

[2]Note that Crowley did not really explain why there is a distinction between the interest in the course and whether they want it; another alternative would be to simply drop the $\mathbf{W_A}$ and $\mathbf{W_B}$ nodes and have each $\mathbf{I_i}$ being the parent of $\mathbf{Happy_i}$ utility node.

(a)

(b)

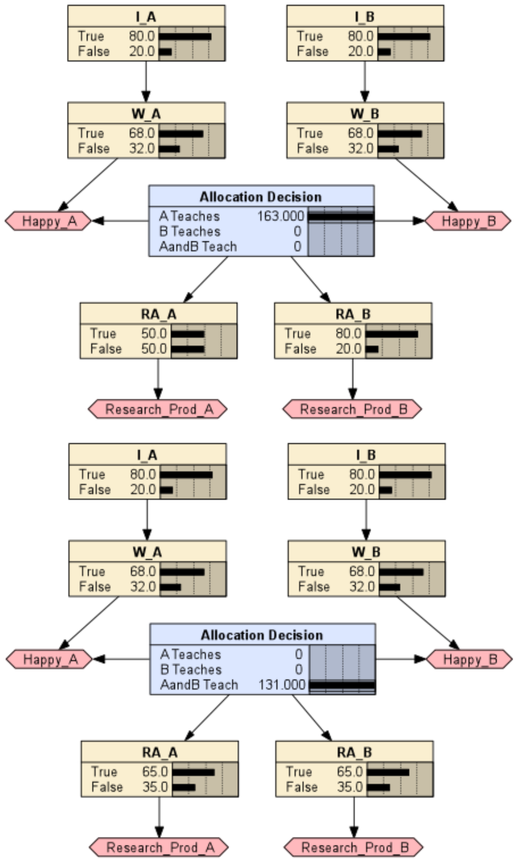

Figure 9: Effect of decision being made (a) A teaches (b) Both teach, neither of which impacts on the **I** and **W** nodes.

5. CONCLUSIONS AND RECOMMENDATIONS

We were looking into the general problem of modelling constraints in BNs (Albrecht and Bud (2009)) when we came across Crowley's teaching BN, and his proposed shielding approach for dealing with unwanted side effects, as described in Section 2. Unsatisfied with the problems still remaining with the BN model for such a seemingly simple problem, we developed a number of alternative models. In this paper we have presented these as a case study in modelling choices, particularly around modelling decisions, which we feel may be of some interest to the BN modelling community. These lead to the following more general recommendations to BN modellers.

First, when problems are identified with a model, for example when belief updating results in unexpected or unwanted changes in the distributions for some variables, rather than introducing new nodes or arcs to "fix" the problem, the modeller should carefully examine the modelling choices embedded in the existing model. Second, it is also important to consider whether the aim is to produce (1) a BN that is a probabilistic model of someone's decision making process, or (2) a decision network model to support decision making of the general modelling choices involved.

We note that there are other approaches to modelling the teaching example that we have not consider here, including the use of object-oriented BNs to scale up across larger number of lecturers and courses, and dynamic Bayesian networks to model explicitly the impact over time of teaching allocations to research output. Also, the impact on complexity on some of the modelling alternatives, only need to be taken into consideration for larger decision problems, e.g., when allocating many lecturers across many units. It may be that at larger scales, the problem is better modelled as a non-probabilistic constraint satisfaction problem, rather than as a BN.

REFERENCES

References

Albrecht, D. and A. Bud (2009), Keeping track of pieces in Invisible Chess, *First Annual Conference of the Australasian Network Modelling Conference*, December, Melbourne.

Crowley, M. (2005), Shielding Against Conditioning Side Effects in Graphical Models, Master's thesis, The University of British Columbia.

Howard, R. and J. Matheson (1981), Influence Diagrams, In R. Howard and J. Matheson (Eds.), *Readings in Decision Analysis*, pp. 763–771. Strategic Decisions Group, Menlo Park, CA.

Jensen, F. and T. Nielsen (2007), *Bayesian Networks and Decision Graphs* (2nd ed.). New York; Springer Verlag.

Korb, K. (2011), Intervention & Observation, In *Proceedings of the Third Annual Conference of the Australasian Bayesian Network Modelling Society (ABNMS2011)*, November, Brisbane.

Korb, K., L. Hope, A. Nicholson, and K.Axnick (2004), Varieties of Causal Intervention, In *PRICAI'04 - Proceedings of the 8th Pacific Rim International Conference on Artificial Intelligence*, Auckland, New Zealand, pp. 322–331.

Korb, K. and A. Nicholson (2010), *Bayesian Artificial Intelligence* (2nd ed.). Chapman & Hall/CRC.

Pearl, J. (2000), *Causality: Models, Reasoning and Inference*. Cambridge, MA.