Towards a Method of Building Causal Bayesian Networks for Prognostic Decision Support

Barbaros Yet[1], Zane Perkins[2], William Marsh[1], Norman Fenton[1]

[1]School of Electronic Engineering and Computer Science, Queen Mary, University of London, Mile End Road, London, E1 4NS, UK
[2]Trauma Unit, Royal London Hospital, Barts and The London NHS Trust, London, E1 1BB, UK
barbaros@eecs.qmul.ac.uk

Abstract. We describe a method of building a decision support system for clinicians deciding between interventions, using Bayesian Networks (BNs). Using a case study of the amputation of traumatically injured extremities, we explain why existing prognostic models used as decision aids have not been successful in practice. A central idea is the importance of modeling causal relationships, both so that the model conforms to the clinicians' way of reasoning and so that we can predict the probable effect of the available interventions. Since we cannot always depend on data from controlled trials, we depend instead on 'clinical knowledge' and it is therefore vital that this is elicited rigorously. We propose three stages of knowledge modeling covering the treatment process, the information generated by the process and the causal relationship. These stages lead to a causal Bayesian network, which is used to predict the patient outcome under different treatment options.

Keywords: Bayesian Networks, Causal Models, Clinical Decision Support

1 Introduction

How can a decision-support system assist a clinician deciding between several available treatments (or 'interventions') for a patient? We describe a method of building a decision support system applicable to this problem, based on the use of Bayesian Networks (BNs). Our focus here is on the prediction of the outcome for the patient, given the different treatment options, as if to answer a clinician asking "what is likely to happen to the patient if I do A or B?". Such a prediction is the first step needed to assist a decision maker; the further step from prediction to advice is not considered here.

We have developed the proposed method as part of a project to develop decision support for the treatment of traumatically injured (or 'mangled') extremities, where surgeons must decide whether or not to salvage or amputate the injured limb. We use this case study as a running example to illustrate each stage of the method.

The use of prognostic models in medicine is increasing [1]. Such models make predictions about the course of a disease from one or more predictors. The relationship between the predictors and the outcome does not always need to be causal [2]. On the other hand, when the need is to decide between possible interventions, a causal relationship between the intervention and the outcome is clearly necessary and this is a challenge when, as in our case study, we are depending on data gathered from past cases rather than from a controlled trial.

Randomised controlled trials (RCT) have been the primary way of identifying and measuring causal relations, since randomisation has the potential to reduce the effect of confounding variables. However, it is not straightforward to conduct RCTs for all questions of interest and the cost and time required for generalizable RCTs can be very high. The impracticality of RCTs is especially pertinent for an application such as the treatment of mangled extremity by amputation. Apart from the obvious practical and ethical issues, before an RCT is run some evidence of the potential benefits is needed and this must come from non-experimental sources.

Our proposal is to develop causal BNs based on a combination of expert medical knowledge and observational data. The knowledge is required to identify the causal relations and the data is used for determining the strengths of these relations. Knowledge is captured through a sequence of models describing the treatment process, the information available and a hierarchy of causal relationships.

The remainder of this paper is organised as follows: the case study about mangled extremity is first presented in Section 2, with Section 3 covering existing work on prognostic models and decision support for mangled extremity treatment. Section 4 presents the proposed method for building causal BNs. Conclusions and discussions are given in Section 5.

2 Case Study: Mangled Extremities

2.1 Treatment of Mangled Extremities

Clinicians often have to decide whether to amputate or salvage the extremity during mangled extremity treatment. This decision, with irreversible consequences for the patient, revolves around three possible adverse outcomes, which change in prominence as the treatment progresses.

1. ***Death.*** There is a risk to the patient's life from the injury to the limb. This risk depends on other injuries that may have been sustained at the same time. This risk is most prominent at the first stage of treatment.
2. ***Limb tissue viability.*** If the limb loses its blood supply for too long, its tissues becomes unviable and amputation becomes inevitable. The viability of the limb tissues is evaluated as the extent of the injury is accessed.
3. ***Non-functional limb.*** A salvaged limb may be more or less functional due to the anatomical problems such as loss of muscle compartments or transected nerves. For some patients a prosthetic limb may be preferable to a non-functional or

painful limb; this outcome becomes more prominent when it is clear that limb salvage is possible.

The clinician's concerns about these three treatment outcomes changes as the treatment progresses. The probabilities of the adverse outcomes are both positively and negatively related with each other so it may not be possible to find a decision that minimises all of them. For example, lengthy reconstruction surgeries can salvage patient's limb, but it can also put the patient's life in danger when the patient is physiologically unwell. In later stages of the treatment, following correction of initial physiology, infections of the damaged limb tissues may again threaten patient's life. Finally, the clinicians may decide to amputate the limb if it is not likely to be functional in the long run. Although the choice of treatment is the same, the underlying reasoning changes significantly through different stages of the treatment.

2.2 Experience of the Trauma Unit at the Barts and the London Hospital

The Royal London Hospital (RLH) is an internationally recognised leader in trauma care and trauma research. The trauma unit is the busiest in the United Kingdom treating over 2000 injured patients last year (2010), a quarter of whom were severely injured. The hospital is also the lead for a network of trauma hospitals, the London Trauma System, which provides specialist trauma care for the millions of people living in London and the South-East of England. This trauma system is believed to be the largest of its kind in the world. As a major trauma centre the hospital provides expedient access to the latest technology, treatments and expert trauma clinicians around the clock. Evidence has shown that people who suffer serious injuries need the highest quality specialist care to give them the best chances of survival and recovery.

The most common cause of injury seen at the Royal London Hospital is road traffic collisions followed by stabbings and falls from a height. Nearly half of the trauma patients have an injury to an extremity or the pelvic girdle, and 1% of these patients end up having lower limb amputations. A large multidiscipline team manages those with severe limb injuries. These devastating injuries carry a high mortality and morbidity in a predominantly young population. The multidiscipline approach ensures the best possible outcome for these patients.

2.3 Characteristics of this Decision Problem

We can summarise the characteristics of the limb amputation decision problem as follows:

- The treatment pathway is complex and the decision evolves with the treatment.
- Multiple outcomes need to be considered.
- The information relevant to the decision changes with time.

These characteristics suggest the need for analysis of the information available and modelling of the care pathway before a decision model can be developed.

3 Prognostic Models

3.1 Traditional Prognostic Models

Prognosis is the act of predicting the course of a disease or a medical condition. A prognostic model makes such predictions based on several independent predictors. Typically, the relation of the predictors to the model outcome is analysed by multivariate statistical models or similar approaches [3]. The accepted way of selecting predictors is to adjust the variables and check their effects on the outcome in observational data. If an adjustment of a variable is connected to the outcome with statistical significance, the variable can be called as an independent predictor. The danger is that correlation is confused with causation. For example, grey hair is an independent risk factor for heart disease, however, if two men of the same age but different hair colours are considered, grey hair does not probably increase the heart disease risk [2]. Therefore, the independent predictors are not necessarily causal factors; they are the factors that are correlated with causal factors according to the available data and selected variables. More extreme examples about variable selection can be seen in some scientific studies where electric-razors or owning refrigerators have been identified as risk factors for cancer [4]. Consequently, the independent predictors and their relations to outcome can be completely different between studies. Predictors with different sets of variables can be statistically accurate but high statistical accuracy of a model does not ensure its clinical acceptance [5] and there are now widely accepted arguments against the use of statistical significance tests and their associated p-values [6]. Clinicians demand models that have reasonable and understandable knowledge base aligned with latest clinical guidelines [7, 8].

On the other hand, there is an abundance of domain knowledge about the clinically relevant variables and their causal relations that can be integrated into model building. The main problems of traditional prognostic approaches can be overcome if domain knowledge is used.

3.2 Scoring Systems for Mangled Extremity Treatment

Multiple scoring systems have been developed as decision support models for mangled extremity treatment [9]. All of these models grade a patient's situation according to several injury-related variables. If a patient's score is above the model's threshold value, the model recommends an amputation. However, these scoring systems have not been widely accepted as a decision support tool by clinicians; we consider some reasons for this below.

Firstly, the scoring systems were developed based on observational data with low sample sizes. For example, MESS [10], which is a widely known scoring system, was developed with data on just 26 patients. Consequently, the high predictive results obtained by the authors were not repeated in later independent validation studies that have a higher number of participants (Table 1). Validation of the model was measured by sensitivity, which is the percentage of the amputated limbs that were also predicted to be amputated by the model, and by specificity, which is the percentage of the salvaged limbs that were predicted as such by the model. Sensitivity and specificity results for the other scoring systems were similar as well. Bosse et al.'s multicentre prospective study [11] concluded that the predictive performance of the scoring systems was poor.

Table 1. Validation Studies for MESS

Validation Study	***Participants***	***Sensitivity***	***Specificity***
MESS's developers [10]	26	1	1
Robertson et al.[12]	154	0.43	1
Bonanni et al.[13]	89	0.22	0.53
Durham et al.[14]	51	0.79	0.83
Bosse et al.[11]	556	0.46	0.91
Korompilias et al.[15]	63	0.87	0.71

Secondly, the output of scoring systems was the amputation decision itself. As a result, if there is a discrepancy between the model's recommendations and clinician's decisions, the model does not provide any useful decision support apart from implying that this outcome was the decision that was made in the model's training data. Thirdly, the scoring system's performance cannot be assessed in practice by sensitivity and specificity values since these measures represent the similarity between the models' recommendations and clinicians' decisions. A model can have 100% sensitivity and specificity but there is a possibility that both model and the compared clinicians were wrong.

3.3 Bayesian Networks

Bayesian networks (BNs) are probabilistic graphical models with multiple variables and relevant independence assumptions that are suitable for representing causality. All BNs, on the other hand, are not necessarily causal since the BNs can effectively represent non-causal probabilistic relations as well as the causal ones. BNs have been proposed for a wide range of medical applications [16] including prognosis [17] and prediction of the outcomes of different interventions [18].

Verduijn et al. [17] proposed a method for learning BNs specifically for prognosis from observational data. Their approach has several advantages compared to traditional prognostic models since it can represent the reasoning mechanism among intermediate variables. Moreover, in contrast to regression models the multiple stage nature of prognostic decisions can be implemented in BNs. Although Verduijn et al.'s prognostic BNs [17] are capable of learning more complex relations from observational data; those relations are still not necessarily causal so that making predictions about interventions is not possible. There are several methods for learning parts of causal relations from data [19] but these methods require extensive amount of data which may not be feasible for relatively uncommon medical conditions such as traumatic amputations.

Causal BNs [19] should have a clear relationship to the complex procedural, associational and hierarchical aspects of the clinical knowledge together with the causal relations. Such knowledge is elicited and verified from multiple experts to minimise the biases. However, communicating through the model becomes more difficult with this additional complexity. Moreover, the risk of introducing a semantic mistake to the model increases.

Several knowledge modelling approaches have been proposed to overcome those difficulties in building BN structure. Nadkarni and Shenoy [20] outline a procedure which can be useful for building simpler causal BNs. Laskey and Mahoney [21] propose using systems engineering methods for building larger and more complex BNs. Object-oriented approaches have been proposed as well to assist the building of larger BNs [22, 23]. Laskey and Mahoney [24] propose using network fragments with object-oriented concepts to represent repeatable structures in the problem domain that are meaningful to the experts. Neil et al. [25] use repeatable structures that represent commonly encountered modelling tasks such as modelling the measurements in BNs. A more automated way of building BN with expert knowledge is proposed by Wiegerinck [26] in which constraints on the model are identified by the experts, and the model is modified by minimising a cost-function which shows the model's differences from those constraints. Although their method is primarily used for tuning model parameters, model structure can be modified as well with the help of the cost function and the constraints. Helsper and van der Gaag [27, 28] propose keeping detailed background knowledge for the BN in a separate ontology from which they gather initial BN alternatives. These alternatives are then modified and improved until one of them is satisfactory for the user. Additional expert knowledge, which is not stored in the ontology, could be necessary for these improvements. Moreover, it is not clear if the aims of the BN, relevant decisions and priorities could be analysed with the ontology. These issues should be clearly identified in the knowledge base for a complex multi-stage decision making problem like the mangled extremity treatment. In the following section, we will give some examples about the challenges of BN building which have not been fully solved by the previous knowledge modelling approaches, and introduce a method addressing them.

4 Knowledge Modelling for Causal Bayesian Networks

Since our proposal to use causal BNs depends on the elicitation of knowledge about causal relationships between variables, explicit knowledge modelling is central to our proposed method. In this section, we describe this knowledge modelling, illustrating it with examples from the case study of mangled extremities.

4.1 Method Overview

Our goal is to develop BN models to predict one or more outcome variables, depending on the values of other relevant factors and conditioned on the possible outcomes. The first imperative is therefore to have a clear understanding of all the variables in the model (i.e. *clarity test* [29]), so before constructing the BN we need to capture knowledge about the entities and attributes relevant to the domain. These entities may relate to different stages of the treatment process and some attributes may have changing values. A complete understanding of the data therefore depends on knowledge of the treatment process. Moreover, the predictions needed for decision support may change through the treatment. A model of this process is therefore our starting point.

4.2 Modelling the Treatment Process

Decisions about clinical interventions are usually done in iterative stages until the patient is treated. After making an intervention, clinicians observe the results of the intervention, re-evaluate treatment risks, and select a treatment alternative [30]. Activity diagrams (Fig. 1) can be used to identify the decisions that are important for the clinical problem, and the priorities of these decisions throughout the treatment.

The changing decision priorities in mangled extremities are illustrated by an example about a patient treated by surgeons at RLH following a motor-cycle accident that resulted in severe leg injury and serious bleeding. When the patient arrives at the hospital, his physiology is in a dangerous condition due to bleeding but his limb appears to be anatomically salvageable. A causal BN used for decision-support at this stage will access the physiology-related risk of death, considering the options of a reconstruction operation, and the possibility of salvaging the limb later. Consequently variables of the model will be mainly about physiology, bleeding and limb injury.

The risk of death related to physiology may decrease if the patient is resuscitated for a few days. However, other risks to the patient's life may develop in the following days. These include infections and renal failure resulting from dead or dying tissues. The causal BN used at this stage will still provide decision support about the risk of death and possibility of limb salvage but its predictions will be based these developing pathologies. If the risk of death related to limb injury is also low, the clinicians will evaluate the possibility of anatomical salvage and future functioning of the limb. The causal BN for this stage will be more focused on structure of the injured limb rather than mechanisms related to death.

Modelling such differences between decision making stages could be complex, especially if there are multiple decisions with various priorities and interrelated outcomes like the mangled extremity treatment example. The activity diagrams provide a clear and understandable map of the main decisions making stages and relevant interventions (shown by diamonds and rectangles respectively in Fig. 1). Most of the medical experts are familiar with the format of activity diagrams as many clinical guidelines are published in similar ways. Therefore, it is also a convenient communication medium between the domain expert and the knowledge engineer.

Not all the information may be needed in the BN to predict the outcomes of interest at each stage. The main outcomes and relevant variables for each stage can be identified by using the activity diagram. For example, when we focus on the first decision making stage about the patient's physiology with the domain experts, the main outcome for this stage is identified as the risk of death, and the relevant clinical variables are identified as bleeding, shock, and coagulation. The main causal relations between the outcome and the clinical variables are identified as: the patient enters a lethal state of circulatory shock as a result of bleeding, shock may impair the body's coagulation ability worsening the future course of bleeding and shock status.

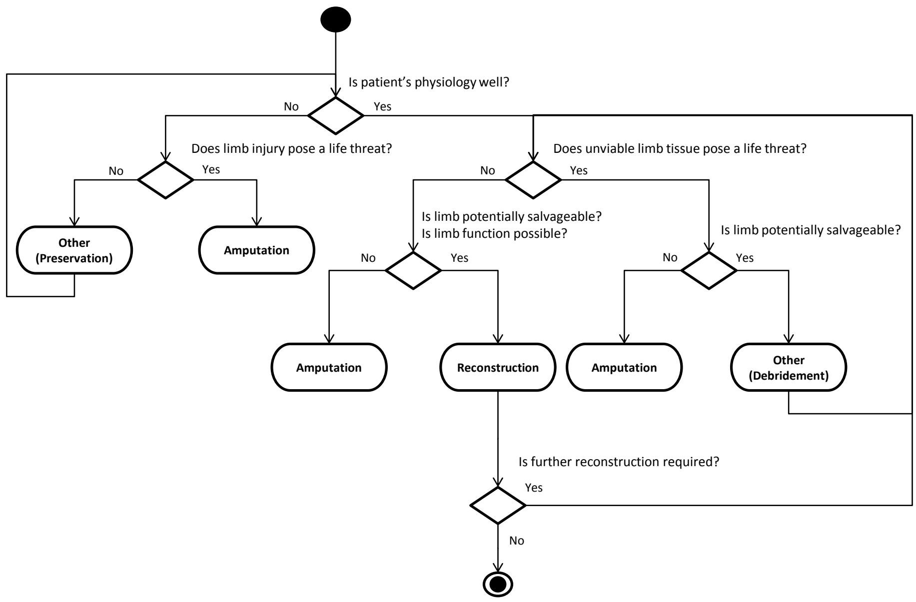

Fig. 1. Activity Diagram for Mangled Extremity Decision Making

4.3 Modelling Information Arising From Treatment

The variables used in the BN must be clearly defined, corresponding to an attribute of a defined entity, at a given stage of treatment. Information models that represent the knowledge about relevant entities and their attributes can guide the selection of variables in the BN. Moreover, multiplicity about these variables must be clarified as well. In our case study, a patient may have an amputation in each of their two limbs. Moreover, the same limb could be sequentially amputated at progressively higher levels. For example, there are records for 53 patients, 73 limbs and 83 amputation operations in the data from RLH about lower limb amputations.

The information model can be used with the activity diagram to identify the variables relevant to each decision making stage. For example, we identified main variables and causal relations for the first decision making stage in Section 4.2 but some of those variables (e.g. shock) are unobservable so that their states must be estimated by other observable attributes (e.g. systolic blood pressure (SBP), heart rate (HR), respiratory rate (RR)). The information model can be used to identify such attributes related to main entities (Fig. 2).

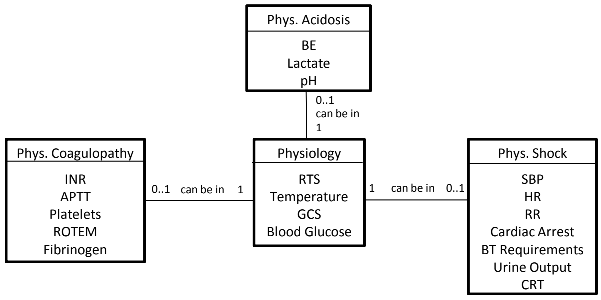

Fig. 2. Fragment of the Information Model about Physiology

Many of the unobservable variables about physiology and their estimators are continuously changing variables but the values for most of these are not measured continuously. For example, multiple blood tests are used to estimate the changes in coagulopathy. Moreover, the causal BN models are used in discrete time stages therefore the relations between the variables with multiple measurements and their representations in the causal BN must be clarified. The class diagrams can be used for illuminating those relations. In the class diagram about mangled extremity treatment (Fig. 3), the model assumes that a patient can have multiple interventions, and the patient's physiology status can change between these interventions. Therefore, the instantiations about a patient's physiology in the causal BN shows the state in each intervention. On the other hand, variables about a patient's past medical history (PMH) or injury are static (Fig. 3) thus they have single fixed values in the causal BN.

4.4 Model Causal Relationships at Different Knowledge Levels

While clinicians usually express their reasoning in small and compact statements, these statements are actually based on series of cause-effect deductions from more complex structures. Methods for representing multiple levels of clinical knowledge have been developed [31]. The causal BNs with less detail abstract the detailed information about a part of a clinical problem. These models can show the main causal relations with fewer variables which is suitable for communication with the experts about the overall model structure. More detailed causal BNs can show more complex relations that could be used for making inferences about detailed mechanisms if there is available data (for example, from a variety of laboratory tests). These models are aligned by the less detailed models through focal nodes. Focal nodes are anchors for the different knowledge levels that describe the same concept and share the same name [31].

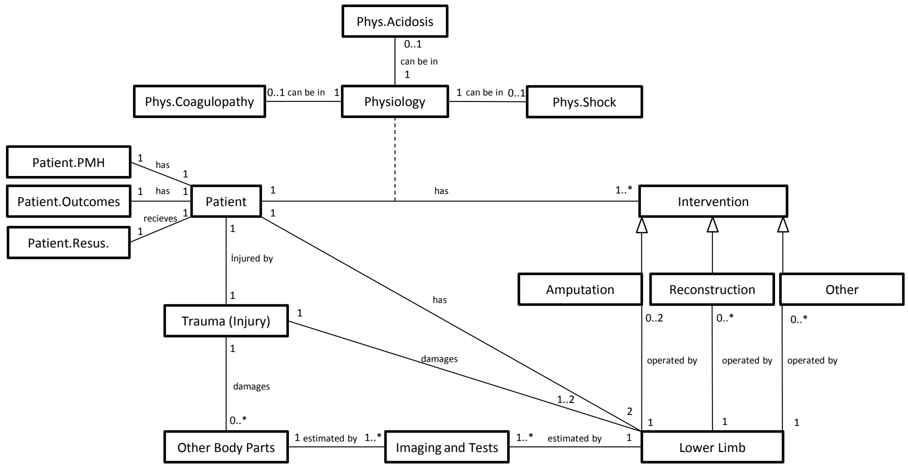

Fig. 3. Class diagram of entities related to mangled extremity treatment

An example of causal BNs with different detail levels is shown in a fragment of the mangled extremity model (Fig. 4). These causal BNs model a part of the physiology related risk of death which is crucial in early stages of the treatment (Fig. 1). The outcomes (death) and main variables (bleeding, shock, coagulopathy) for the model were identified in Section 4.2 with the help of the activity diagram. The feedback relationship between coagulopathy and future course of bleeding has not been represented as a dynamic BN in this illustration for simplicity. The least detailed causal BN shows the overall causal relations between bleeding, circulatory shock, coagulopathy, the risk of death and possible interventions i.e. amputation or rapid surgery. Although this model represents the overall causal relationships, it does not show the two intermediate (temperature, acidosis) variables between shock and coagulopathy. A more detailed version of the causal BN can be built by adding these relations as well as the estimators for the unobservable shock variable (RR, HR, SBP, capillary refill time (CRT), urine output, Glasgow coma scale (GCS)) which were identified by the information model in Section 4.3 (Fig. 2). This model could bring more explanatory predictions due to additional causal mechanisms. The relation between shock and its seven estimators can also be explained in a more detailed way. For example, urine output that is used for estimating shock is caused by perfusion in the kidneys. The increase in respiratory rate is caused by lack of O_2 delivery to the tissues as a result of low perfusion. Therefore, knowledge detail in the model can be increased by modelling shock through these relations. However, estimating values about the perfusion in different body parts could be more difficult for the user than estimating a value for shock only. The nodes that are not modelled in different levels of details, such as bleeding or coagulopathy node in our example can be used as focal points to align the models and keep the overall causal relations consistent between different detail levels.

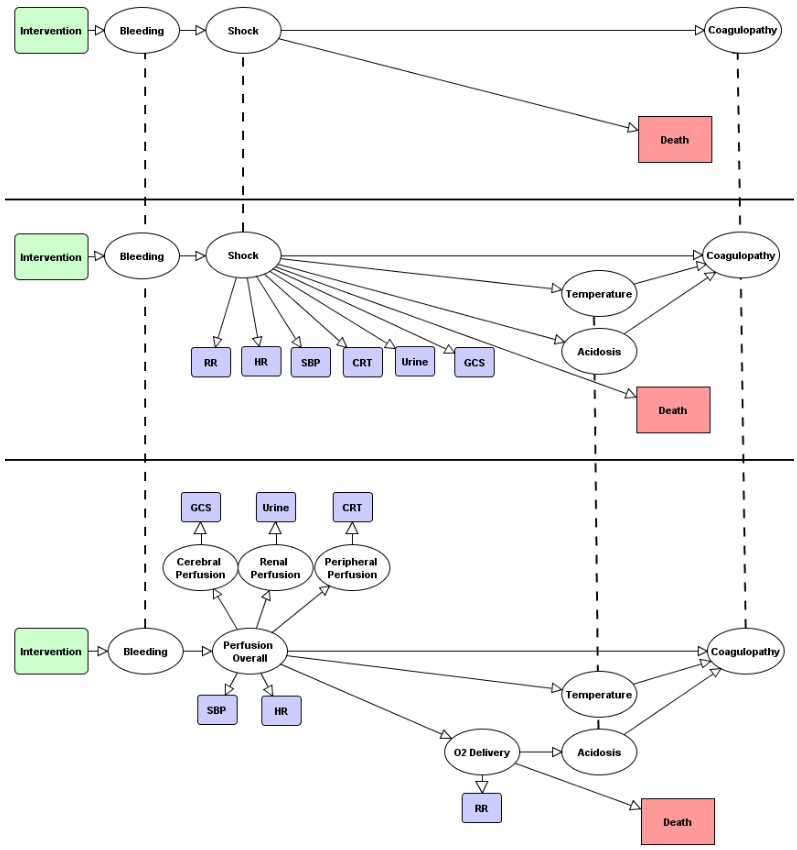

Fig. 4. Causal BN with multiple levels of detail about physiology related risk of death

Any of the three models in the example can be used depending on the available information about the variables, predictive performance of the model and preference of the user. The multi-level causal BN makes it possible to keep a consistent and understandable knowledge-base for the model regardless of the modelling preferences. This could be useful in improving the model's clinical acceptance since a clear and reliable knowledge-base is one of the main demands from the prognostic models [7, 8].

4.5 Modelling Dynamic Variables

Many continuous clinical variables are estimated by multiple discrete measurements such as blood tests. The multiplicity relations between these variables and their measurements can be identified by the information models shown in Section 4.3. On the other hand, modelling the effects of continuously changing variables in the BN still remains an issue. One well-known solution for this issue is to instantiate the complete model structure over multiple time slices. However, this approach could be computationally infeasible if there are numerous time stages and large model structures.

One approach for modelling continuously changing variables in the BN could be to use trend variables that summarise the variations of several previous instantiations of the related variable. In clinical practice, the trends of historical measurements for some clinical factors are used to make predictions about patient outcomes. For example, a patient's response to resuscitation, which can be analysed by trends of several diagnostic values about shock, is an important factor for predicting the patient's survival. This can be modelled in BN by adding a trend variable that summarises the variations in the previous states of shock (Fig 5.).

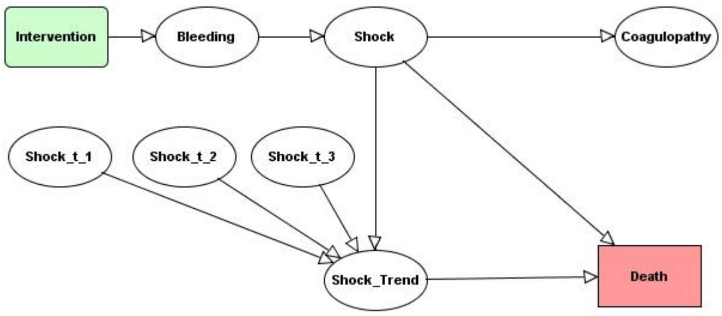

Fig. 5. Fragment of the causal BN with a trend variable

5 Conclusion

In this study, we have proposed a method for building causal BNs, where causal relationships are elicited from clinical knowledge. The method involves three stages of knowledge modelling, using:

- activity diagrams to model the decision points and procedural relations
- class diagrams to model the multiplicity relations between the variables
- multi-level causal diagrams to represent a hierarchical of causal relationships.

This method aids the knowledge-elicitation with experts by providing understandable intermediate models and decreases the risk of having semantic mistakes in the final BN model. The study for developing the method is still in progress. This paper shows our first attempts for providing guideline for some common modelling problems seen in building causal BNs. More structured method for building complete causal BNs are being researched. For next steps, we plan to formalise the models within a common framework, allowing more automated approaches for building the final causal BNs. The outcomes of the causal BN are posterior probability distributions about the treatment risks in a variety of situations. Although these posterior distributions can provide useful information for the decision maker, we plan to analyse these distributions' relations to decision making and prepare clinical guidelines that are more helpful and efficient for the decision maker.

Acknowledgements

We are grateful for the contribution for Mr Nigel Tai, FRCS, Consultant Trauma and Vascular Surgeon at the Barts and the London NHS Trust to the work described in this paper.

References

1. Moons, K., Royston, P., Vergouwe, Y., Grobbee, D., Altman, D.: Prognosis and prognostic research: what, why, and how? BMJ 338, 1317--1320 (2009)
2. Brotman, D.J., Walker, E., Lauer, M.S., O'Brien, R.G.: In Search of Fewer Independent Risk Factors. Arch. Intern. Med. 165, 138--145 (2005)
3. Abu-Hanna, A., Lucas, P.J.F.: Prognostic Models in Medicine. Method. Inform. Med. 40, 1--5. (2001)
4. Jenks, S., Volkers, N.: Razors and refrigerators and reindeers – oh my. J. Natl. Cancer. Inst. 84, 1863 (1992)
5. Moons, K.G.M., Altman, D.G., Vergouwe, Y., Royston, P.: Prognosis and prognostic research: application and impact of prognostic models in clinical medicine. BMJ 338, 1487--1490 (2009)
6. Ziliack, S.T., McCloskey, D.N.: The cult of statistical significance : How the standard error costs us job, justice, and lives. University of Michigan Press, Ann Arbor (2008)
7. Wyatt, C.J., Altman, D.G.: Commentary: prognostic models: clinically useful or quickly forgotten? BMJ 311, 1539--1541 (1995)
8. Akhthar, J.I., Forse, R.A.: Prognostic models: Are these models health fortune-telling tools? Crit. Care. Med. 38(7), 1605--1606 (2010)
9. Rush, R.M. Jr., Beekley, A.C., Puttler, E.G., Kjorstad, R.J.: The Mangled Extremity. Curr. Probl. Surg. 46(11), 851--926 (2009)
10. Johansen, K., Daines, M., Howey, T., Helfet, D., Hansen, S.T.Jr.: Objective criteria accurately predict amputation following lower extremity trauma. J. Trauma. 30, 568--572 (1990)
11. Bosse, M.J., MacKenzie, E.J., Kellam, J.F., Burgess, A.R., Webb, L.X., Swiontkowski, M.F., Sanders, R.W., Jones, A.L., McAndrew, M.P., Patterson, B.M., McCarthy, M.L., Cyril, J.K.: A Prospective Evaluation of the Clinical Utility of the Lower-Extremity Injury-Severity Scores. J. Bone. Joint. Surg. 83-A(1), 3--14 (2001)
12. Robertson, P.A.: Prediction of Amputation after Severe Lower Limb Trauma. J. Bone. Joint. Surg. 73-B(5), 816--818. (1991)
13. Bonanni, F., Rhodes, M., Lucke, J.F.: The Futility of Predictive Scoring of Mangled Lower Extremities. J. Trauma. 34, 99--104 (1993)
14. Durham, R.M., Mistry, B.M., Mazuski, J.E., Shapiro, M., Jacobs, D.: Outcome and Utility of Scoring Systems in the Management of the Mangled Extremity. Am. J. Surg. 172, 569--574 (1996)

15. Korompilias, A.V., Beris, A.E., Lykissas, M.G., Vekris, M.D., Kontogeorgakos, VA, Soucacos, PN: The mangled extremity and attempt for limb salvage. J. Orthop. Surg. Res. 4(4). (2009)
16. Lucas, P.J.F, van der Gaag, L.C., Abu-Hanna, A.: Bayesian networks in biomedicine and health-care. Artif. Intell. Med. 30(3), 201--214 (2004)
17. Verduijn, M., Peek, N., Rosseel, P.M.J., de Jonge, E., de Mol, B.A.J.M.: Prognostic Bayesian Networks I: Rationale, learning procedure and clinical use. J. Biomed. Inform. 40(6), 609--618 (2007)
18. Fenton, N., Neil, M.: Comparing risks of alternative medical diagnosis using Bayesian arguments. J. Biomed. Inform. 43, 485--495 (2010)
19. Pearl, J.: Causality. Cambridge University Press, Cambridge. (2000)
20. Nadkarni, S., Shenoy, P.P: A Bayesian network approach to making inferences in causal maps. Eur. J. Oper. Res. 128, 479--498 (2001)
21. Laskey, K., Mahoney, S.: Network engineering for agile belief networks. IEEE T. Knowl. Data. En. 12, 487--498 (2000)
22. Koller, D., Pfeffer, A.: Object-Oriented Bayesian Networks. In: Proceedings of the Thirteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-97), pp. 302--313. (1997)
23. Bangsø, O., Wuillemin, P.H.: Top-down Construction and Repetitive Structures Representation in Bayesian Networks. In: Proceedings of the Thirteenth International Florida Artificial Intelligence Research Society Conference, pp. 282--286. (2000)
24. Laskey, K., Mahoney, S.: Network Fragments: Representing Knowledge for Constructing Probabilistic Models. In: Proceedings of the Thirteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-97), pp. 334--341.
25. Neil, M., Fenton, N., Nielsen, L.: Building large-scale Bayesian networks. Knowl. Eng. Rev. 15(3), 287--284. (2000)
26. Wiegerinck, W.: Modeling Bayesian Networks by Learning from Experts. In: Proceedings of the 17th Belgium-Netherlands Conference on Artificial Intelligence (BNAIC 2005), pp. 305--310. (2005)
27. Helsper, E.M., van der Gaag, L.C.: Building Bayesian networks through ontologies. In: 15th Eurpoean Conference on Artificial Intelligence, pp. 680--684. (2002)
28. Helsper, E.M., van der Gaag, L.C.: Ontologies for probabilistic networks: a case study in the oesophageal-cancer domain. Knowl. Eng. Rev. 22(1), 67--86. (2007)
29. Kjærulff, U.B., Madsen, A.L.: Bayesian Networks and Influence Diagrams: A Guide to Construction and Analysis. Springer, New York. (2008)
30. Sox, H.C., Marshal, A.B., Michael, C.H., Keith, I.M.: Medical Decision Making. ACP, Philadelphia.(1988)
31. Patil, R.S., Szolovitz, P., Schwartz, W.B.: Causal Understanding of Patient Illness in Medical Diagnosis. In: Seventh International Joint Conference on Artificial Intelligence, pp.893--899. (1981)