

Combining data and meta-analysis to build Bayesian networks for clinical decision support

Barbaros Yet^{1,*}, Zane B. Perkins², Todd E. Rasmussen³, Nigel R. Tai⁴, and D. William R. Marsh¹

¹School of Electronic Engineering and Computer Science, Queen Mary University of London, UK

²Centre for Trauma Science, Queen Mary University of London, UK

³The United States Army Medical Research and Materiel Command, Maryland, USA

⁴The Royal London Hospital, London, UK

*Corresponding Author: Barbaros Yet

Risk and Information Management Research Group, Room CS332, School of Electronic Engineering and Computer Science, West Square, Queen Mary University of London, E1 4NS, London, UK.

E-mail: b.yet@qmul.ac.uk

Tel: +44 20 7882 8027

Abstract:

Complex clinical decisions require the decision maker to evaluate multiple factors that may interact with each other. Many clinical studies, however, report ‘univariate’ relations between a single factor and outcome. Such univariate statistics are often insufficient to provide useful support for complex clinical decisions even when they are pooled using meta-analysis. More useful decision support could be provided by evidence-based models that take the interaction between factors into account. In this paper, we propose a method of integrating the univariate results of a meta-analysis with a clinical dataset and expert knowledge to construct multivariate Bayesian network (BN) models. The technique reduces the size of the dataset needed to learn the parameters of a model of a given complexity. Supplementing the data with the meta-analysis results avoids the need to either simplify the model –

1 ignoring some complexities of the problem – or to gather more data. The method is illustrated by a
 2 clinical case study into the prediction of the viability of severely injured lower extremities. The case
 3 study illustrates the advantages of integrating combined evidence into BN development: the BN
 4 developed using our method outperformed four different data-driven structure learning methods, and a
 5 well-known scoring model (MESS) in this domain.

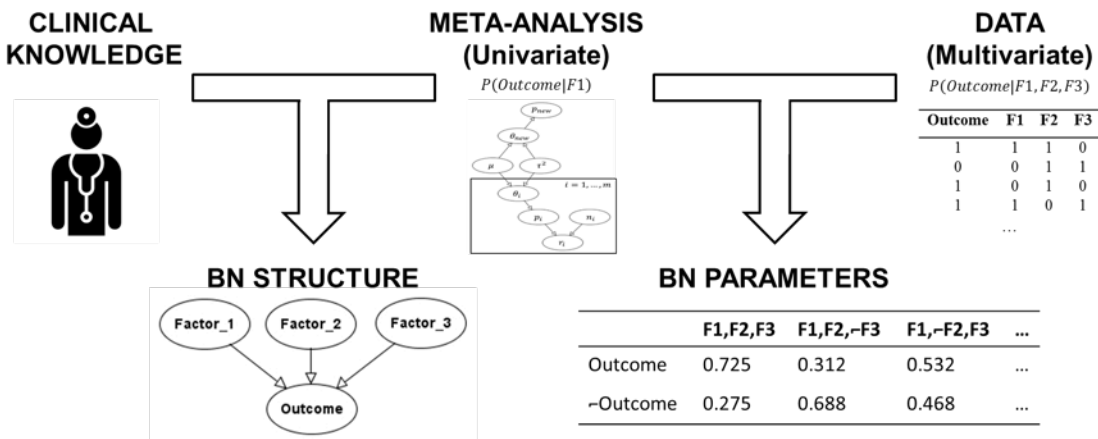
6 **Keywords:**

7 Clinical decision support; Bayesian networks; evidence-based medicine; evidence synthesis; meta-
 8 analysis

9 **Highlights:**

- 10 • We focus on complex clinical problems where data is available in small amounts.
- 11 • Our methodology compensates for the lack of data by using published evidence.
- 12 • It combines multivariate data with univariate meta-analysis to build BN models.
- 13 • The method is illustrated by a medical case study on trauma care.
- 14 • Our method outperformed data-driven techniques and MESS model in the case study.

15 **Graphical Abstract:**



16

17

1

2 1 Introduction

3 It is a challenge to build effective decision-support models for complex clinical problems; such
4 problems involve multiple interacting factors [1,2] and to account for both the factors and their
5 interaction a ‘multivariate’ model is needed [3]; these can have many forms: our focus is on Bayesian
6 networks. In general, a multivariate model can be built in a number of ways: 1) purely from data using
7 statistical and machine learning techniques [4], 2) from a combination of clinical knowledge and data
8 [5–7] or 3) from published literature using multivariate meta-analysis techniques [8]. Each of these
9 techniques has been shown to be successful in certain conditions but in this paper, we focus on
10 clinical problems where none of these techniques is sufficient, on its own, to build a useful decision
11 support model. That is, our focus is on problems that are complex, important but also rare: their rarity
12 makes it hard to collect very large datasets (so called ‘big data’); their complexity demands a
13 sophisticated multivariate model but their importance ensures that a large number of relevant research
14 studies is available.

15 In these domains, the first method of building models – purely from data – results in simple models
16 that cannot deal with the complexity of the problem [1] because there is not enough data to support a
17 complex model. The third approach fails because clinical studies rarely publish information detailed
18 enough for multivariate meta-analysis [9]. Instead, many medical studies report ‘univariate’ relations
19 between a single factor and an outcome. Randomised controlled trials, for example, analyse the effect
20 of a single treatment by using randomisation to decrease the confounding effect of other variables
21 [10]. Similarly, many observational studies report the relation between individual risk factors and
22 outcomes even when their dataset contains information about multiple factors. The second approach –
23 combining knowledge and data – could work but it ignores the large body of published evidence; our
24 challenge is therefore to exploit the results of a meta-analysis of studies reporting univariate relations
25 to supplement a dataset that is otherwise inadequate to support a complex multivariate model.

1 Decision support directly from univariate relations is limited, as the effects of interactions between
2 variables are not taken into account. For example, evidence about individual effects of a treatment and
3 a comorbidity factor can be analysed in separate meta-analyses. However, if the treatment and
4 comorbidity factor interact with each other, their joint effect may be completely different from their
5 individual effects. As a result, decision support provided by the meta-analysis of individual effects
6 may be invalid for a patient who is exposed to both the treatment and the comorbidity factor (see
7 [2,10,11] for a more detailed discussion of generalising clinical evidence).

8 To improve this situation, we propose a method of combining the results of meta-analyses, clinical
9 knowledge and data to provide decision support for complex decision problems where the data is
10 scarce. Our method combines ‘univariate’ meta-analysis following a systematic review, with a small
11 ‘multivariate’ dataset and expert knowledge. Bayesian networks (BN) offer a powerful framework to
12 combine evidence from different sources [1,5,12,13]. Our methodology integrates the evidence from a
13 meta-analysis into BN development by using it first to identify the BN structure and then to help
14 determine the BN parameters; this second step uses auxiliary parameter learning models similar in
15 some ways to techniques that can be used for meta-analysis. We illustrate the application and results
16 of this method with a clinical case study into the prediction of the outcomes of severely injured lower
17 extremities.


18 In the remainder of this paper, Section 2 recaps of a Bayesian meta-analysis technique to combine
19 probabilities. Section 3 describes our methodology for developing a BN based on the results of a
20 meta-analysis. Sections 4, 5 and 6 present the case-study, results and conclusions respectively.

21 **2 Meta-analysis of Probabilities**

22 The method we propose in Section 3 assumes a possibly small multivariate patient dataset is available
23 together with univariate results of a meta-analysis of probabilities. In this section, we give a recap of
24 the meta-analysis of probabilities by briefly presenting an existing Bayesian technique [14,15]. The
25 results obtained from this meta-analysis technique can be used in the method of Section 3, though

1 other techniques could also be used. The recap also serves to introduce hierarchical Bayesian models,
 2 which are also used in Section 3.

3 Meta-analysis is an important form of clinical evidence as it combines and summarises the relevant
 4 published evidence that is identified by a systematic literature review. Meta-analysis can be used to
 5 combine different types of statistics including odds ratios, risk ratios and probabilities [14]. We focus
 6 on the meta-analysis of probabilities as the parameters of a BN are composed of probabilities. Figure
 7 1 shows a random-effects Bayesian meta-analysis model that takes the variation between studies into
 8 account, and does not assume normality for the distribution of the individual studies.



9

10 *Figure 1 Bayesian meta-analysis model for combining probabilities*

11 The binomial distribution is the probability distribution of the number of positive outcomes in n
 12 independent experiments where the probability of a positive outcome is p for every experiment. In the
 13 meta-analysis model, the result of each individual study i is modelled with the binomial distribution
 14 shown below, where r_i is the number of positive outcomes observed in the study i , p_i is the true study
 15 probability of the study i , and n_i is the sample size of the study i .

$$r_i \sim \text{Binomial}(p_i, n_i)$$

16

1 The normal distribution is a convenient way of modelling the pooled estimate and the variation
2 between studies. We use an inverse logit transformation to model the true study probability p_i with
3 the normal distribution. The mean μ of this distribution represents the transformed pooled estimate,
4 and the variance τ^2 represents the variation between studies.

$$\text{logit}(p_i) = \theta_i$$

5

$$\theta_i \sim \text{Normal}(\mu, \tau^2)$$

6

7 The predictive probability distribution can also be calculated by using an inverse logit transformation
8 of this normal distribution. The predictive distribution is a recommended way of presenting the results
9 of a meta-analysis as it represents the uncertainty from both the pooled estimate and the variation
10 between studies (see [14] and chapter 8 of [16] for more detailed information on predictive
11 distributions in meta-analysis) .

$$\theta_{new} \sim \text{Normal}(\mu, \tau^2)$$

12

$$\text{logit}(p_{new}) = \theta_{new}$$

13

14 Finally, priors must be selected for the pooled estimate and between-study standard deviation. The
15 non-informative priors shown below can be used if informative priors are not available.

$$\mu \sim \text{Normal}(0, 1000)$$

16

$$\tau \sim \text{Uniform}(0, 2)$$

17

18 In order to calculate the posteriors of μ , τ^2 and p_{new} , we enter the observed number of positive
19 outcomes r_i and sample sizes n_i from each reviewed study. The posteriors can be calculated by using

1 the dynamic discretisation algorithm [17] in AgenaRisk [18] or the Markov Chain Monte Carlo
2 (MCMC) sampling technique in OpenBUGS [19].

3 **3 Building BNs based on Meta-analysis**

4 The previous section described a Bayesian meta-analysis technique for pooling probabilities. In this
5 section, we present a methodology that uses data, expert knowledge and the pooled probabilities from
6 a meta-analysis to define the structure (Section 3.1) and parameters (Section 3.2) of a BN decision
7 support model. Our methodology assumes that expert knowledge, a meta-analysis of univariate
8 relations from a relevant systematic review and some data about multivariate relations are available.
9 However, the amount of data may be insufficient to learn the parameters of some relations in the BN.

10 **3.1 Structure**

11 A BN structure can be developed in two stages: selecting variables, and identifying the relations
12 between those variables. Domain experts use the results of a meta-analysis to select the important
13 variables for the BN. The experts review every variable that is considered to be clinically important in
14 the meta-analysis. During the review, they define mechanistic relations between each of these
15 variables and the outcome. These definitions enable us to 1) build a causal BN structure that is
16 consistent with clinical knowledge 2) identify the variables that are clinically important considering
17 the aims and scope of the model. Variables that are outside the scope of the model are excluded even
18 when they have a clinically significant effect in the meta-analysis.


19 The mechanistic relations between the observed factors and outcome may depend on clinical factors
20 that are not available in the data or not examined in the meta-analysis. For example, the data may not
21 distinguish between the measurements and true state of a variable, or may exclude a part of the
22 important causal factors in the domain (see [6] and chapters 1 and 2 of [20] for a more detailed
23 discussion of this issue). In this case, latent variables are used to model clinical knowledge in the BN.

1 3.2 Parameters

2 Meta-analysis of a univariate relation provides a probability conditioned on a single variable, such as
3 $P(Y|X_1)$. Such probability distribution cannot be directly used for a BN variable that is conditioned
4 on multiple parents such as $P(Y|X_1, \dots, X_n)$. In this section, we present a parameter learning method
5 for combining the results of a univariate meta-analysis and data to learn the parameters of a BN
6 variable that has multiple parents. Our method uses auxiliary Bayesian models to learn the parameters
7 of the BN used for decision support. These auxiliary models are hierarchical models with a structure
8 that is conceptually similar to the Bayesian meta-analysis model described in Section 2. We introduce
9 the proposed method by a simple example in Section 3.2.1, and examine the application of this
10 method to more complex BN models in Section 3.2.2. The conditional probabilities provided by a
11 meta-analysis may be relevant to variables that are not directly linked in the BN structure. We
12 examine this issue in Section 3.2.3.

13 3.2.1 Illustration of the Parameter Learning Method

14 In this section, we introduce our parameter learning method with the simple BN shown in Figure 2.



15

16 *Figure 2 Simple BN used for illustrating the parameter learning method*

17 This BN has 3 variables and each of its variables has 2 states:

18
$$X_1 = \{x_1^1, x_1^2\}$$

19
$$X_2 = \{x_2^1, x_2^2\}$$

20
$$Y = \{y^1, y^2\}$$

1 The conditional probability distribution (CPD) of a discrete variable is encoded in a node probability
 2 table (NPT) in a BN. Table 1 shows the NPT of the variable Y . We require 4 parameters for this NPT:
 3 $P(y^1|x_1^1, x_2^1)$, $P(y^1|x_1^1, x_2^2)$, $P(y^1|x_1^2, x_2^1)$ and $P(y^1|x_1^2, x_2^2)$.

4 *Table 1 NPT of the variable Y*

	$X_1 = x_1^1$	$X_1 = x_1^1$	$X_1 = x_1^2$	$X_1 = x_1^2$
	$X_2 = x_2^1$	$X_2 = x_2^2$	$X_2 = x_2^1$	$X_2 = x_2^2$
$Y = y^1$	$P(y^1 x_1^1, x_2^1)$	$P(y^1 x_1^1, x_2^2)$	$P(y^1 x_1^2, x_2^1)$	$P(y^1 x_1^2, x_2^2)$
$Y = y^2$	$1 - P(y^1 x_1^1, x_2^1)$	$1 - P(y^1 x_1^1, x_2^2)$	$1 - P(y^1 x_1^2, x_2^1)$	$1 - P(y^1 x_1^2, x_2^2)$

5 The parameters of the variable Y can be learnt from data using the maximum likelihood estimate
 6 (MLE) approach. For example, $P(y^1|x_1^1, x_2^1)$ can be estimated by dividing $M[y, x_1^1, x_2^1]$ to $M[x_1^1, x_2^1]$,
 7 where $M[y, x_1^1, x_2^1]$ represents the count of data instances where $Y = y^1$, $X_1 = x_1^1$ and $X_2 = x_2^1$, and
 8 $M[x_1^1, x_2^1]$ represents the count of data instances where $X_1 = x_1^1$ and $X_2 = x_2^1$.

$$P(y^1|x_1^1, x_2^1) = \frac{M[y^1, x_1^1, x_2^1]}{M[x_1^1, x_2^1]}$$

9
 10 Suppose we have a dataset, with a sample size of $M = 250$, to learn the parameters of the BN in
 11 Figure 2. Figure 3 shows a part of the relevant counts from this imaginary dataset. There are only 3
 12 data instances where $Y = y^1$, $X_1 = x_1^1$ and $X_2 = x_2^1$ as shown by $M[y^1, x_1^1, x_2^1] = 3$ in this figure.

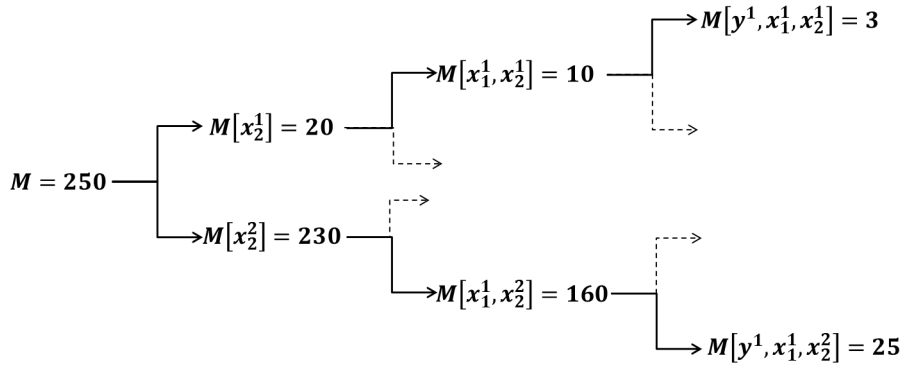


Figure 3 Relevant counts from the data

Our aim is to estimate the parameters of the BN. Although the overall sample size of the data is not small, there is not an adequate amount of data for learning some of the parameters. For example, there are only a few data instances to learn the probability of $P(y^1|x_1^1, x_2^1)$ since $M[y^1, x_1^1, x_2^1] = 3$ and $M[x_1^1, x_2^1] = 10$.

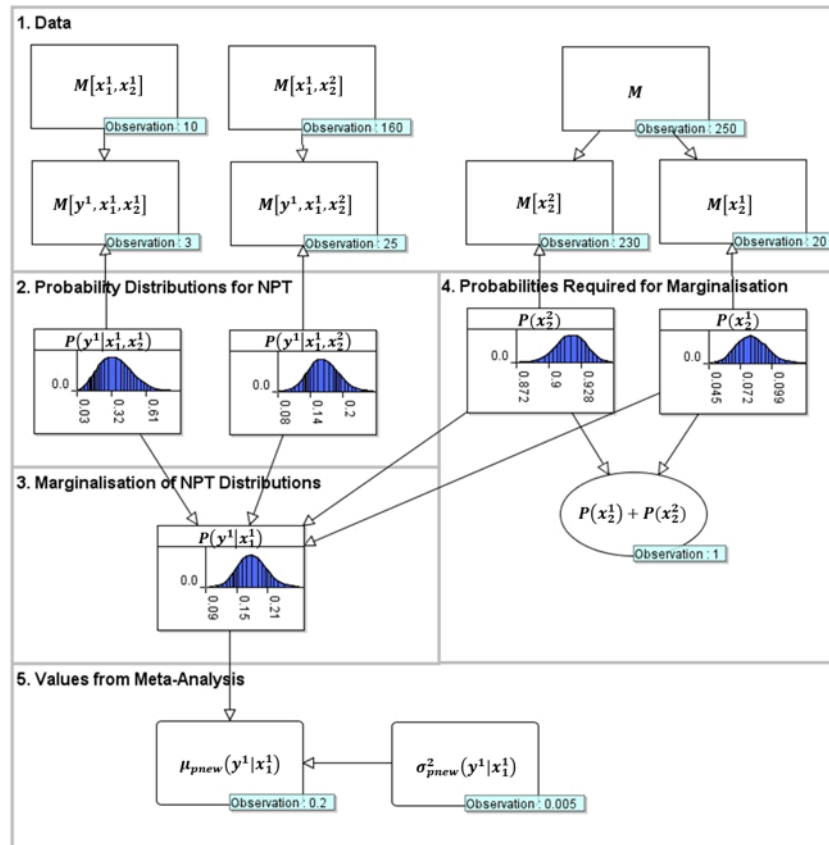
As well as the data, suppose we have the results of a meta-analysis that analyses the relation between Y and X_1 . This meta-analysis pools the conditional probabilities of $P(y^1|x_1^1)$ reported in different studies. The result of the meta-analysis is reported by the mean, $\mu_{pnew}(y^1|x_1^1)$, and variance, $\sigma_{pnew}^2(y^1|x_1^1)$, of the predictive distribution of the pooled conditional probability (see Table 2). A way of calculating these statistics is described in Section 2.

Table 2 Predictive distribution parameters from the meta-analysis

Meta-analysis of $P(y^1 x_1^1)$	
Predictive Distribution Parameters	
$\mu_{pnew}(y^1 x_1^1)$	0.2
$\sigma_{pnew}^2(y^1 x_1^1)$	0.005

The results of the meta-analysis cannot be directly used for the BN parameters since the variable Y is conditioned on both X_1 and X_2 in the BN model whereas it is conditioned only on X_1 in the meta-

1 analysis. In other words, there is no parameter to use $P(y^1|x_1^1)$ directly in the NPT of the variable Y
 2 (see Table 1).
 3 In the remainder of this section, we present a novel technique that combines the data shown in Figure
 4 3 and the meta-analysis results shown in Table 2 to learn the parameters $P(y^1|x_1^1, x_2^1)$ and
 5 $P(y^1|x_1^1, x_2^2)$ for the NPT of the variable Y . The generalisation of this method for a larger number of
 6 parents and states is described in Section 3.2.2.



7
 8 *Figure 4 BN Representation of the Auxiliary Parameter Learning Model*

9 Figure 4 shows a BN representation of the implemented technique. The BN representation is divided
 10 into five components that are described in the remainder of this section:

- 11 1. **Data:** This part uses the binomial distribution to model the relation between the conditional
 12 probability distributions (CPD) that we aim to estimate and the observed counts in the data.

1 For example, the number of data instances where $X_1 = x_1^1, X_2 = x_2^1$ and $Y = y^2$, shown by
 2 $M[y^1, x_1^1, x_2^1]$, has a binomial distribution where the probability parameter is $P(y^1|x_1^1, x_2^1)$
 3 and the number of trials parameter is $M[x_1^1, x_2^1]$. The binomial distributions used in this part
 4 are shown below:

$$M[y^1, x_1^1, x_2^1] \sim \text{Binomial}(M[x_1^1, x_2^1], P(y^1|x_1^1, x_2^1))$$

$$M[y^1, x_1^1, x_2^2] \sim \text{Binomial}(M[x_1^1, x_2^2], P(y^1|x_1^1, x_2^2))$$

$$M[x_2^1] \sim \text{Binomial}(M, P(x_2^1))$$

$$M[x_2^2] \sim \text{Binomial}(M, P(x_2^2))$$

9 2. **Probability Distributions for NPT:** This part contains the CPDs that we aim to estimate for
 10 the NPT of Y . We assign uniform priors for these distributions, informative expert priors can
 11 also be used when available:

$$P(y^1|x_1^1, x_2^1) \sim \text{Uniform}(0,1)$$

$$P(y^1|x_1^1, x_2^2) \sim \text{Uniform}(0,1)$$

13 3. **Marginalisation of NPT Distributions:** Since the variable Y is conditioned only on 1
 14 variable in the meta-analysis and 2 variables in the BN, we model the probability distribution
 15 from the meta-analysis, $P(y^1|x_1^1)$, as the marginalisation of the probability distribution from
 16 the BN parameters $P(y^1|x_1^1, x_2^1)$ and $P(y^1|x_1^1, x_2^2)$:

$$P(y^1|x_1^1) = \sum_{X_2} (P(y^1|x_1^1, X_2) * P(X_2)) = P(y^1|x_1^1, x_2^1)P(x_2^1) + P(y^1|x_1^1, x_2^2)P(x_2^2)$$

18

1 4. **Probabilities Required for Marginalisation:** In order to calculate the marginalisation in part
 2 3, we need the probability distributions of $P(x_2^1)$ and $P(x_2^2)$. We assign uniform priors for
 3 these variables. We also assign a constraint to ensure that the sum of $P(x_2^1)$ and $P(x_2^2)$ is
 4 equal to 1.

$$P(x_2^1) \sim \text{Uniform}(0,1)$$

$$P(x_2^2) \sim \text{Uniform}(0,1)$$

$$\sum_{x_2} P(X_2) = P(x_2^1) + P(x_2^2) = 1$$

5
 6
 7
 8 5. **Values from Meta-analysis:** The pooled estimate $\mu_{pnew}(y^1|x_1^1)$ from the meta-analysis is
 9 modelled with the normal distribution truncated to a unit interval as it represents a probability
 10 value, denoted by $TNormal_{[0,1]}(\mu, \sigma^2)$. We use $P(y^1|x_1^1)$ from the marginalisation in part 3
 11 and $\sigma_{pnew}^2(y^1|x_1^1)$ from the predictive distribution as the mean and variance of this normal
 12 distribution respectively. The values from the meta-analysis are modelled as:

$$\mu_{pnew}(y^1|x_1^1) \sim TNormal_{[0,1]}(P(y^1|x_1^1), \sigma_{pnew}^2(y^1|x_1^1))$$

13
 14 After the observations from the data and meta-analysis are entered to the BN (see Figure 4), the
 15 posteriors for $P(y^1|x_1^1, x_2^1)$ and $P(y^1|x_1^1, x_2^2)$ can be calculated. Note that, the NPT of Y requires
 16 point estimates for $P(y^1|x_1^1, x_2^1)$ and $P(y^1|x_1^1, x_2^2)$ whereas our model calculates the entire probability
 17 distribution of these parameters. Therefore, we take the mean of these distributions for the point
 18 estimates required for the NPT (see Section 17.4 of [21] for a discussion of the use of posterior
 19 distributions for BN parameters).


1 In the following section, we describe the generalisation of this technique for estimating the parameters
 2 of variables with more parents or states.

3 3.2.2 Application of the Parameter Learning Method for More Complex BNs

4 Let Y be a BN variable that has n parents, and $\mathbf{X} = \{X_1, X_2, \dots, X_n\}$ be the set of parents of Y (see
 5 Figure 5). Both Y and its parents have multiple states:

6
$$Y = \{y^1, \dots, y^k\}$$

7
$$X_i = \{x_i^1, \dots, x_i^k\}$$



8

9 *Figure 5 BN with n parents used for illustrating the generalised parameter learning method*

10 Our dataset contains a total of M data instances about \mathbf{X} and Y (see Table 3). We also have pooled
 11 conditional probability and variance estimates of the predictive distribution of $P(Y|X_i)$ from a meta-
 12 analysis (see Table 4). The way of calculating these predictive distributions is described in Section 2.

13 *Table 3 Sample learning dataset*

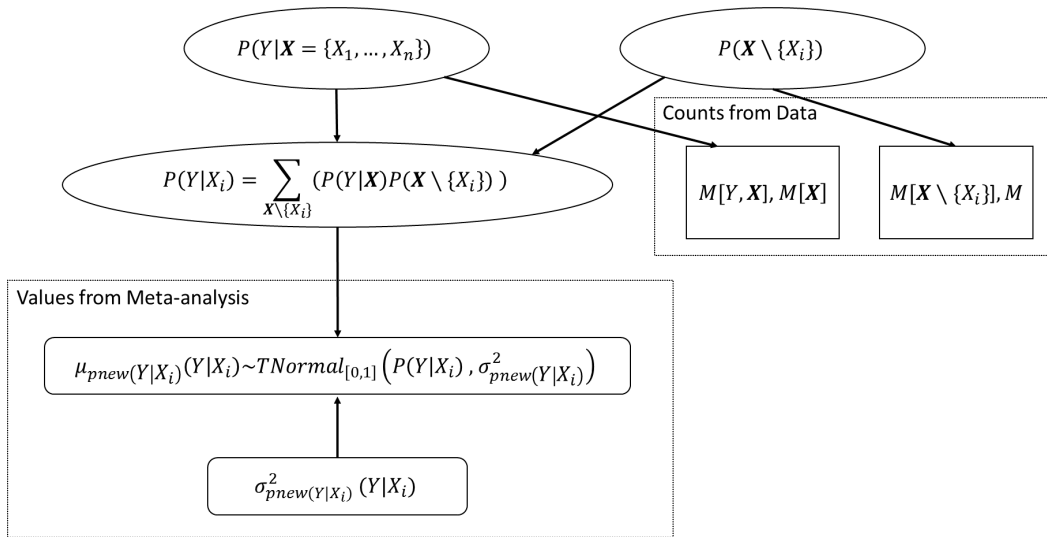
	Y	X₁	...	X_n
1	y^4	x_1^3	\dots	x_2^2
2	y^2	x_1^2	\dots	x_2^1
:	:	:		:
:	:	:		:
M	y^1	x_1^1	\dots	x_2^4

14 Figure 6 shows an abstract graphical illustration of the generalised auxiliary parameter learning
 15 model. This model is a generalisation of the model shown in Figure 4. This illustration is not a BN; it

1 is a schema for building an auxiliary parameter learning model for any number of states and parent
 2 variables. The size of the auxiliary parameter learning model grows rapidly with increasing number of
 3 parents and states.

4 *Table 4 Sample meta-analysis results*

	$\mu_{P_{new}}$	$\sigma_{P_{new}}^2$
$P(\mathbf{Y} \mathbf{X}_1)$.	0.13	0.007
$P(\mathbf{Y} \mathbf{X}_2)$.	0.21	0.025
:	:	:
:	:	:
$P(\mathbf{Y} \mathbf{X}_n)$.	0.19	0.001



5
 6 *Figure 6 Graphical Illustration of the Generalised Auxiliary Parameter Learning Model*

7 In Figure 6, the variables shown by ellipses are unknown variables that will be estimated by the
 8 model. The variables shown by rounded rectangles are observed with the values from the meta-
 9 analysis, and the variables shown by rectangles are observed from the dataset. The constraints that
 10 sum probabilities to 1 are not included in this figure to simplify the illustration. By running this
 11 auxiliary model, we estimate probability distributions for the parameters $P(\mathbf{Y}|\mathbf{X})$ required by the NPT
 12 of Y . Since the BN requires only a point estimate of the parameter, not the entire distribution; we use
 13 the mean of this distribution as the BN parameter.

1 According to our model, the data related to Y , i.e. $M[Y, \mathbf{X}]$, is generated by the binomial distribution
 2 with the probability of success $P(Y|\mathbf{X})$ and the number of trials $M[\mathbf{X}]$.

$$M[Y, \mathbf{X}] \sim \text{Binomial}(M[\mathbf{X}], P(Y|\mathbf{X}))$$

3
 4 $M[Y, \mathbf{X}]$ represents the count of data instances for specific values of X_1, \dots, X_n and Y . For example,
 5 $M[y^2, x_1^1, x_2^3, \dots, x_n^4]$ represents the number of data instances where
 6 $Y = y^2, X_1 = x_1^1, X_2 = x_2^3, \dots, X_n = x_n^4$. Similarly $M[\mathbf{X}]$ represent the number of data instances where
 7 X_1, \dots, X_n have certain values.

8 Our aim is to estimate the CPD of $P(Y|\mathbf{X})$. We assign a uniform prior for this distribution;
 9 informative expert priors can also be used when available.

$$P(Y|\mathbf{X}) \sim \text{Uniform}(0,1)$$

10
 11 The meta-analysis results are conditioned on a fewer variables than the CPD in the BN. Therefore, the
 12 expected values of the meta-analysis results are modelled as a marginalisation of the CPD. The meta-
 13 analysis provided the pooled conditional probability estimates about $P(Y|X_i)$ that are modelled as the
 14 marginalisation of $P(Y|\mathbf{X})$

$$P(Y|X_i) = \sum_{\mathbf{X} \setminus \{X_i\}} P(Y|\mathbf{X})P(\mathbf{X} \setminus \{X_i\})$$

15

1 $P(\mathbf{X} \setminus \{X_i\})$ is also estimated by the binomial distribution below where M denotes the total number of
 2 data instances, and $M[\mathbf{X} \setminus \{X_i\}]$ denotes the counts of data instances with $\mathbf{X} \setminus \{X_i\}$. $P(\mathbf{X} \setminus \{X_i\})$ has a
 3 uniform prior

$$M[\mathbf{X} \setminus \{X_i\}] \sim \text{Binomial}(M, P(\mathbf{X} \setminus \{X_i\}))$$

$$P(\mathbf{X} \setminus \{X_i\}) \sim \text{Uniform}(0,1)$$

4
 5
 6 The pooled estimates from the meta-analysis $\mu_{P_{new}(Y|X_i)}$ are modelled with a normal distribution
 7 truncated to a unit interval [0–1] as it represents a probability. The mean of this distribution is the
 8 marginalisation of the CPD, i.e. $P(Y|X_i)$, and the variance $\sigma_{P(Y|X_i)}^2$ represents the degree of
 9 uncertainty we assign to the meta-analysis results. We enter the mean and variance of the predictive
 10 distribution in meta-analysis as observations for $\mu_{P_{new}(Y|X_i)}$ and $\sigma_{P_{new}(Y|X_i)}^2$. We use the truncated
 11 normal distribution as it is convenient to define the expected value and variance parameters for it but
 12 $\mu_{P_{new}(Y|X_i)}$ may not be normally distributed as it represents a probability value between 0 and 1.

$$\mu_{P_{new}(Y|X_i)} \sim \text{TNormal}_{[0,1]}(P(Y|X_i), \sigma_{P_{new}(Y|X_i)}^2)$$

13
 14 Finally, we introduce constraints to ensure that the sum of every probability distribution is equal to 1.

$$\sum_Y P(Y|\mathbf{X}) = 1$$

$$\sum_{\mathbf{X} \setminus \{X_i\}} P(\mathbf{X} \setminus \{X_i\}) = 1$$

$$\sum_Y P(Y|X_i) = 1$$

1

2 3.2.3 Meta-analysis results for non-neighbour variables

3 The method described in Sections 3.2.1 and 3.2.2 assumes that the variables analysed in the meta-
 4 analysis are neighbours in the BN. In this section we look at how this assumption can be relaxed to
 5 handle the more general case where the BN contains other – intermediate – variables between the
 6 variables analysed in the meta-analysis. This situation is illustrated in Figure 7; the meta-analysis
 7 combines the published probabilities for $P(Y|X)$ but the BN contains another variable I between X
 8 and Y so that the values in $P(Y|X)$ are no longer parameters of the BN. We examine how we can use
 9 information from a meta-analysis about non-neighbouring variables to estimate the parameters of a
 10 variable in the BN. In particular, we use $P(Y|X)$, calculated from a meta-analysis, to estimate $P(I|X)$,
 11 giving the parameters of the $X \rightarrow I$ relation, when we know or have data for $P(Y|I)$ describing the
 12 other intermediate relation $I \rightarrow Y$.



13

14 *Figure 7 BN with an intermediate variable*

15 Since every variable in a BN is conditioned on its parents, $P(Y|X)$ provided from the meta-analysis is
 16 equal to:

$$P(Y|X) = \sum_I P(Y|I)P(I|X)$$

17

18 Based on this, we can estimate every parameter of $P(I|X)$ as:

$$P(i^k|X) = \frac{P(Y|X) - \sum_{m \in S} P(Y|i^m)P(i^m|X)}{P(Y|i^k)}$$

1

2 In this equation, S is the set of states of the variable I except the state i^k :

$$S = Val(I) \setminus \{i^k\}$$

3

4 Consequently, the parameters of $P(I|X)$ can be estimated given that meta-analysis provides us with

5 $P(Y|X)$, and we know or have data to learn $P(Y|I)$. In order to get a point estimate for parameters of

6 $P(I|X)$, the number of the states of I must not exceed the number of the states of Y . Otherwise, we

7 can get an interval for the values in $P(I|X)$ but cannot estimate the exact values.

8 For example, let X, I and Y in Figure 7 have two states each with the values $\{x^1, x^2\}, \{i^1, i^2\}, \{y^1, y^2\}$

9 respectively. Suppose a meta-analysis provides us with $P(y^1|x^1) = 0.8$, and we learn the

10 probabilities $P(y^1|i^1) = 0.9$, $P(y^1|i^2) = 0.3$ from the data. Since I has two states:

11 $P(i^2|x^1) = 1 - P(i^1|x^1)$. From these values, we can calculate $P(i^1|x^1)$ as:

$$P(i^1|x^1) = \frac{P(y^1|x^1) - P(y^1|i^2)P(i^2|x^1)}{P(y^1|i^1)} = \frac{P(y^1|x^1) - P(y^1|i^2)(1 - P(i^1|x^1))}{P(y^1|i^1)} = 0.833$$

12

13 When I has more states than Y , we cannot get the exact values of the BN parameters but we can find

14 an interval of possible values. Let I has the states $\{i^1, i^2, i^3\}$ instead of $\{i^1, i^2\}$. Since I has three

15 states, $P(i^3|x^1) = 1 - P(i^1|x^1) - P(i^2|x^1)$. Let $P(y^1|i^3) = 0.6$, and the other values be the same

1 as in the example above. The parameters of $P(i^1|x^1)$ can be calculated by solving the equation below,

2 for each independent value of Y :

$$P(i^1|x^1) = \frac{P(Y|x^1) - (P(Y|i^2)P(i^2|x^1) + P(Y|i^3)P(i^3|x^1))}{P(Y|i^1)}$$

3

4 In this case, we cannot get an exact value for $P(I|x^1)$ as the intermediate variable I has more states

5 than Y , resulting in more unknowns than equations. Instead, $P(I|x^1)$ can get any value as long as it

6 satisfies the following conditions:

$$P(i^2|x^1) = 1.5 - 2P(i^1|x^1)$$

7

$$P(i^3|x^1) = -0.5 + P(i^1|x^1)$$

8

$$0 \leq P(I|x^1) \leq 1$$

9

10 We could use expert knowledge, by eliciting additional constraints, to narrow down the set of


11 acceptable values for the parameters of I . In our example, $P(i^1|x^1)$ can get any value between 0.5 and

12 0.75 to satisfy the conditions above. However, some of these values may not make sense to the

13 domain experts, and we can eliminate these values by adding additional constraints. For example, the

14 experts could say that $P(i^1|x^1)$ should only get values above 0.7 and we could reflect it by adding the

15 $P(i^1|x^1) > 0.7$ constraint to the conditions above.



16

17

Figure 8 BN with Multiple Intermediate Variables

1 The technique described above can also be applied when more intermediate variables are present.
 2 Figure 8 shows a BN that has n intermediate variables between Y and X . A similar case, where 2
 3 intermediate variables are present, is encountered in the case study described in the following section
 4 (see Section 4.4.1). An exact estimate can be found given that we know or have data for $P(Y|X)$,
 5 $P(Y|I_k)$ for $k = 2, \dots, n$, and that Y does not have fewer states than I_1 .

$$P(i_1^k|X) = \frac{P(Y|X) - \sum_{m \in S} P(Y|I_n)P(I_n|I_{n-1}) \dots P(I_2|i_1^m)P(i_1^m|X)}{P(Y|i_1^k)}$$

6

7

$$\text{where } S = \text{Val}(I_1) \setminus \{i_1^k\}$$

8 4 Case-study

9 Using the method described in Section 3, we developed a BN to predict viability of a Lower
 10 Extremity with Vascular Trauma (LEVT). This section presents the development of the LEVT BN,
 11 and Section 5 presents its results.

12 4.1 Background

13 Injuries to the blood vessels of the lower extremity are potentially devastating and can result in death,
 14 severe disability or limb loss. Delays or errors in treatment decisions may lead to irreversible
 15 consequences and worsen outcome. One of the most difficult surgical decisions is whether to attempt
 16 salvage or perform an amputation of a severely injured extremity. Accurate risk stratification and
 17 outcome prediction, for a given injury pattern, has the potential to improve outcome by reducing
 18 delays and errors in decision-making.

19 Limb tissues may be permanently damaged as a direct consequence of the energy transfer during
 20 injury or die because of a prolonged disruption to their blood supply. The extent of tissue damage, or
 21 loss, is directly related to future outcome and is the primary determinant of the need for amputation.
 22 Following a lower extremity vascular injury, early reperfusion of the affected tissues is essential to

1 ensure their viability. Reperfusion entails surgical reconstruction of the damaged blood vessels.
2 Predicting the outcome of vascular reconstruction and the projected tissue viability would inform
3 treatment decisions and risks. We developed a BN model for predicting the viability of a traumatic
4 lower extremity with vascular injury after salvage is attempted. The BN is built in collaboration with
5 the Trauma Sciences Unit at the Royal London Hospital and the United States Army Institute of
6 Surgical Research (USAISR). Two trauma surgeons (the 2nd and 4th authors) were involved in
7 development of the LEVT BN. A dataset of 521 lower extremity injuries and 487 patients collected by
8 USAISR, and a systematic review and meta-analysis of the relevant prognostic factors were used to
9 develop the LEVT BN.

10 4.2 Meta-analysis for Lower Extremity Vascular Trauma

11 A number of research studies that describe the factors that affect outcome following LEVT have been
12 published. Our first step was to conduct a systematic review of these studies and perform a meta-
13 analysis of the factors. The systematic review included 45 articles containing information regarding
14 3164 lower extremity repairs. The study protocol is published in the PROSPERO register of
15 systematic reviews [22].

1

Table 5 Mean and Variances of the Predictive Distributions from the Meta-analysis for $P(\text{Amputation}|\text{Clinical Factor})$

Clinical Factor	Predictive Distribution	
	$\mu_{P_{new}}$	$\sigma_{P_{new}}^2$
Arterial Repair		
Graft	0.11	0.009
Primary Repair	0.05	0.002
Anatomical Site		
Femoral	0.04	0.004
Popliteal	0.14	0.005
Tibial	0.10	0.018
Associated Injuries		
MAI* - present	0.18	0.045
MAI* - absent	0.09	0.006
Soft tissue - present	0.26	0.066
Soft tissue - absent	0.08	0.009
Fracture - present	0.14	0.013
Fracture - absent	0.02	0.001
Nerve - present	0.12	0.022
Nerve - absent	0.06	0.016
Complications		
Shock - present	0.13	0.047
Shock - absent	0.06	0.030
Ischaemia time > 6 hrs.	0.24	0.050
Ischaemia time \leq 6 hrs.	0.05	0.009
CS ⁺ - present	0.28	0.008
CS ⁺ - absent	0.06	0.002

*MAI= Arterial Injuries at Multiple Levels, ⁺CS: Compartment Syndrome

2 We used the model described in Section 2 to pool the relevant conditional probabilities and calculate
3 the predictive distributions. The meta-analysis models were calculated using AgenaRisk [18]. Table 5
4 shows the means and variances of the posterior predictive distributions. In the following sections, we
5 use these results to define the structure and parameters of a BN model.

6 4.3 Deriving the BN structure

7 The structure of the BN was defined using the methodology described in Section 3.1. The systematic
8 review identified clinical factors that are potentially associated with the outcome of interest. A meta-
9 analysis of these factors identified the strength of this association (see Table 5). A domain expert (2nd
10 author) examined these variables and described the mechanistic relation between each of the variables
11 and the outcome. These relations were modelled in a causal BN structure. Knowledge of mechanistic
12 relations enabled us to identify variables outside the intended scope of the BN. For example, nerve

1 injuries were not included in the model even though they are associated with an increased probability
 2 of amputation. The domain expert identified that although nerve injuries may affect function of the
 3 related tissues (sensation and movement) they do not affect viability of the tissue. As the intended
 4 scope of our model was to predict limb viability, this variable was excluded.

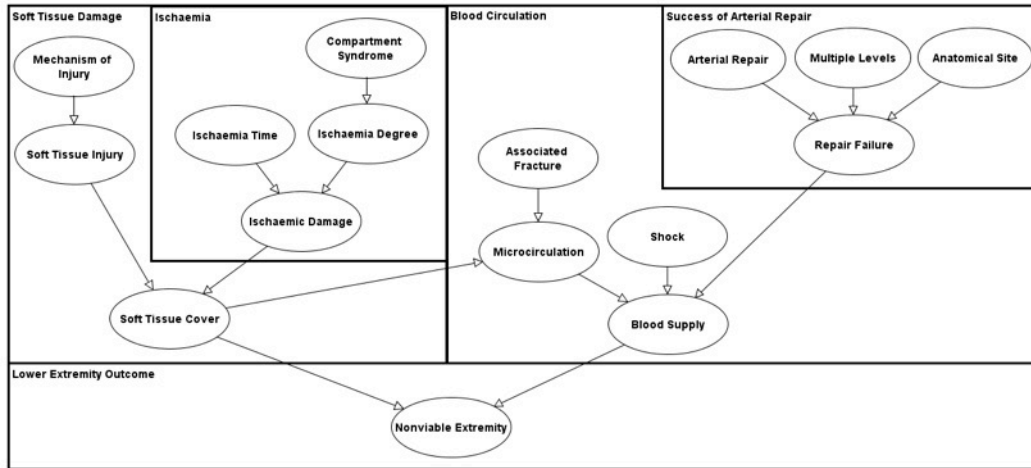
5 *Table 6 Observed and Latent Variables in LEVT BN*

Observed Variables	Latent Variables
Arterial Repair	Blood Supply
Anatomical Site	Ischaemic Damage
Multiple Levels (MAI)	Microcirculation
Soft Tissue Injury	Soft Tissue Cover
Associated Fracture	
Shock	
Ischaemia Time	
Ischaemia Degree	
Compartment Syndrome	
Repair Failure	
Number of Injured Tibials	
Nonviable Extremity	

6 For some variables, our dataset held more detailed information than the results of the meta-analysis.
 7 For example, soft-tissue injury was identified as one of the most important prognostic variables by the
 8 meta-analysis, the information in our dataset allowed us to model this variable in more detailed states
 9 for the BN. Similarly, detailed information on the degree of ischaemia was present in the dataset but
 10 not in the meta-analysis. Therefore, the BN models some variables in more detail than the information
 11 obtained from the meta-analysis.

12 Several latent variables were introduced as the domain expert identified the mechanistic relations
 13 between the observed clinical factors and outcomes. These variables were clinically important but
 14 neither the dataset nor the reviewed studies contained them as they cannot be directly observed [6].
 15 For example, both soft tissue injuries and vessel injuries that require a graft repair have high
 16 probabilities of amputation in the meta-analysis. However, each of these factors is related to
 17 amputation through a different mechanism. Graft repairs can lead to amputation if the graft blocks or
 18 bursts, and thus disrupts the blood flow to the extremity. A variable representing the degree of ‘blood
 19 supply’ is required to model this relation. Although the degree of blood supply can be estimated by
 20 several measurements, the precise state of this variable is difficult to observe and therefore is not
 21 available in the dataset. Soft tissue injuries lead to amputation if insufficient tissue remains to allow

1 repair. Similarly, a latent variable representing the degree of ‘soft tissue cover’ is required to model
 2 this relation. Table 6 shows a list of the observed and latent variables in the LEVT BN structure.
 3 These variables and the LEVT BN structure are described in the following section.



4
 5 *Figure 9 Lower Extremity Vascular Trauma Bayesian Network*

6 The LEVT BN is divided into 5 components, corresponding to the 5 boxes shown in Figure 9. A
 7 summary of the variables and relations in each of these components are shown below:

- 8 • **Lower Extremity Outcome:** A viable lower extremity requires an adequate blood supply and
 9 sufficient viable soft tissue to allow a repair. The ‘Nonviable Extremity’ variable represents
 10 extremities that are amputated as a result of insufficient viable tissue. ‘Nonviable Extremity’
 11 is the main outcome variable that the LEVT BN aims to predict.
- 12 • **Ischaemia:** Tissue ischaemia results when there is an imbalance between the supply of
 13 oxygen to tissue and the tissues oxygen requirements to sustain life. This results from a
 14 disruption in the blood supply to the tissue. Initially ischaemia may be reversible, but if
 15 prolonged will result in permanent death of the affected tissues. Since our model is built for
 16 lower extremities with vascular injuries, most of the extremities within the scope of our model
 17 will be partly or completely ischaemic until the vascular injury is repaired. The severity of
 18 ischaemic damage depends on the time elapsed since the beginning of ischaemia (Ischaemia
 19 Time) and the degree of obstruction (Ischaemia Degree). A second important cause of
 20 ischaemia is the development of a complication called compartment syndrome. Compartment

1 syndrome results when the swelling of injured tissues compresses the blood vessels,
2 disrupting the blood supply.

3 • **Soft Tissue Damage:** This part of the model predicts the projected amount of viable soft
4 tissue in the lower extremity. A critical amount of soft tissue is necessary to repair the lower
5 extremity and protect it from infection. Therefore, the degree of soft tissue cover is one of the
6 main factors affecting limb viability. Our model estimates the amount of soft tissue (Soft
7 Tissue Cover) based on the amount of non-viable tissue due to the direct damage from the
8 injury (Soft Tissue Injury) and ischaemia (Ischaemic Damage). Certain injury types
9 (Mechanism of Injury), such as blast injuries, are likely to cause more severe soft tissue
10 injuries.

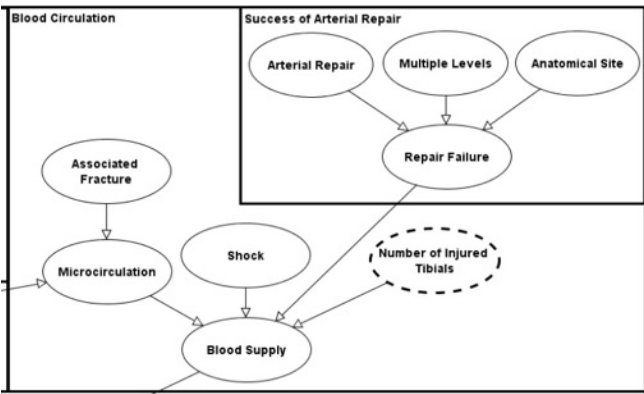
11 • **Success of Arterial Repair:** This part of the model predicts the success of a vascular repair
12 operation represented by the 'Repair Failure' variable. 'Arterial Repair' variable represents
13 the type of the repair operation, and have two states: 'Graft' and 'Primary Repair'. 'Graft'
14 represents bypassing of the injured artery by a vein harvested from the patient. 'Primary
15 repair' represents a simpler repair operation such as stitching of a small laceration in the
16 artery. 'Graft' repairs have higher rate of failure compared to 'Primary Repair' as this
17 operation is more complex and applied to more severe cases. Injury characteristics often
18 define the type of the arterial repair. For example, an arterial injury cannot be treated by
19 primary repair if a significant part of the artery is missing, thus a graft is necessary.

20 The 'Multiple Levels' variable represent whether vascular injuries are present at multiple
21 levels of the same extremity. Repairs of such injuries have a higher probability of failure as
22 they are more likely to block.

23 'Anatomical Site' variable represents the location of the main arterial injury. Our model
24 includes injuries above the knee (femoral artery), at the knee (popliteal artery) or below the
25 knee (tibial arteries). Reconstruction of a femoral artery often has better outcomes compared
26 to a popliteal or a tibial artery.

1 • **Blood Circulation:** ‘Blood Supply’ variable represents the degree of blood supply to the
 2 lower extremity. This variable essentially depends on the ‘Repair Failure’ variable. If the
 3 vascular repair fails, the extremity will not have adequate blood supply; so there is a
 4 deterministic relation between the negative repair failure and inadequate blood supply. In
 5 other words, a repair failure leads to inadequate blood supply, and inadequate blood supply
 6 leads to a non-viable extremity in our model. However, a successful arterial repair may not
 7 guarantee adequate blood supply throughout the lower extremity; side factors including
 8 ‘Shock’ and ‘Microcirculation’ can also affect the outcomes. The ‘Shock’ variable represents
 9 an overall deficiency of blood supply throughout the body. The ‘Microcirculation’ variable
 10 represents the severity of injury in the smaller vessels of the lower extremity.

11 A single vessel supplies blood to the lower extremity. This divides into three branches, called tibial
 12 arteries, below the knee. Modelling tibial arteries is important since, in this segment, limb viability is
 13 related to the number of tibial arteries injured. In order to model this difference, we modified the BN
 14 structure for injuries below the knee by adding a variable about the number of injured tibial arteries.
 15 This modification is shown by the variable with dashed lines in Figure 10. Our model assumes that a
 16 repair failure leads to a non-viable extremity if all 3 tibial arteries are injured. However, there is a
 17 chance of a successful outcome if only 1 or 2 tibial arteries are injured. Apart from this difference,
 18 the BN models for above the knee, at the knee and below the knee injuries are exactly the same.



19
 20

Figure 10 LEVT BN fragment for Injuries Below the Knee

1 4.4 Learning Parameters

2 After building the BN structure, we examined the amount of data available for learning each
 3 parameter in the BN. The data were insufficient, or not available, to learn the parameters of some
 4 relations. In this section, we describe the techniques used to learn the parameters when they have 1)
 5 insufficient amount of data, 2) no data at all (latent variables), and 3) adequate amount of data.

6 4.4.1 Variables with Insufficient Data

7 When there is insufficient data to learn a parameter, we combined the results of the meta-analysis and
 8 data by using the technique described in Section 3.2. We used the mean and variance of the relevant
 9 predictive distributions from the meta-analysis (see Table 5) as observations to the $\mu_{P_{new}}$ and $\sigma_{P_{new}}^2$
 10 variables in the auxiliary learning model of this technique (see Figure 6). For example, Table 7 shows
 11 the amount of available data to learn the NPT of the ‘repair failure’ variable. The data is not evenly
 12 distributed to learn this NPT. While there is plenty of data to learn some parameters, only a few
 13 instances is available to learn others. We used the meta-analysis results for the parameters that have
 14 less than 20 instances of data (shown by bold fonts in 7). We observed that using meta-analysis results
 15 did not change the probabilities significantly when it was applied to the parameters with more than 20
 16 instances of data in our case study. Therefore, we used a purely data driven approach for those
 17 parameters (see Section 4.4.3). However, it should be noted that different thresholds may be suitable
 18 to other applications in different domains.

19 *Table 7 Amount of Data Available for Learning Parameters of Repair Failure Variable*

AR*	Graft	Graft	Graft	Graft	Graft	Graft	Primary	Primary	
MAI*	True	True	True	False	False	False	True	True	...
AS*	Femoral	Popliteal	Tibial	Femoral	Popliteal	Tibial	Femoral	Popliteal	
RF*									
Number of observations:	14	6	2	71	115	38	1	3	...

*AR: Arterial Repair, MAI: Multiple Levels, AS: Anatomical Site, RF: Repair Failure

20 The meta-analysis provides us the pooled probabilities of an unsuccessful outcome conditioned on
 21 each individual clinical factor (see Table 5). The variable equivalent to an unsuccessful outcome is
 22 ‘nonviable extremity’ in the LEVT BN but the meta-analysis results can also be used for the NPT of

1 the ‘repair failure’ variable as 1) our model assumes a deterministic relation between an unsuccessful
 2 outcome and repair failure, and thus we know the parameters of the intermediate variables between
 3 them 2) the parents of ‘repair failure’ can influence ‘nonviable extremity’ through only one pathway
 4 (see Section 3.2.3). For example, the ‘arterial repair’ variable can affect ‘nonviable extremity’ through
 5 the following pathway in our model:

Arterial Repair → *Repair Failure* → *Blood Supply* → *Nonviable Extremity*

(AR) (RF) (BS) (NE)

6
 7
 8 The probability provided by the meta-analysis, $P(NE = True|AR = Graft)$, is equivalent to
 9 marginalisation of ‘repair failure’ and ‘blood supply’ from this pathway:

$$P(NE = True|AR = Graft) = \sum_{BS,RF} P(NE = True|BS)P(BS|RF)P(RF|AR = Graft)$$

10

11 In our model, a repair failure always leads to inadequate blood supply, and inadequate blood supply
 12 always leads to a nonviable extremity so that:

$$P(BS = Low|RF = True) = 1, P(NE = True|BS = Low) = 1,$$

$$P(BS = Low|RF = \neg True) = 0, P(NE = True|BS = \neg Low) = 0$$

14

15 By using these values in the marginalisation equation above, we get

16 $P(NE = True|AR = Graft) = P(RF = True|AR = Graft)$. Consequently, we can use the

17 probabilities from the meta-analysis for learning the relation between ‘repair failure’ and its parents.

18 *Table 8 Parameters learnt purely from data and from a combination of data and meta-analysis*

AR*:	Graft	Graft	Graft		Primary	Primary	
MAI*:	True	True	True	...	True	True	...
AS*:	Femoral	Popliteal	Tibial		Femoral	Popliteal	
RF*:						

True	0.17	<i>0.17</i>	0.39	<i>0.41</i>	0.99	<i>0.49</i>		0.01	<i>0.14</i>	0.01	<i>0.29</i>	...
False	0.83	<i>0.83</i>	0.61	<i>0.59</i>	0.01	<i>0.51</i>		0.99	<i>0.86</i>	0.99	<i>0.71</i>	
Number of observations:	14		6		2			1		3		

**AR: Arterial Repair, MAI: Multiple Levels, AS: Anatomical Site, RF: Repair Failure*

1 We used the OpenBUGS software [19] to calculate the posteriors of the auxiliary learning model. We
2 used the Gelman and Rubin diagnostic technique [23], sample plots and autocorrelation plots to assess
3 the convergence. We discarded the first 10,000 samples in MCMC as the burn-in samples, and
4 calculated the posterior distributions based on the next 70,000 samples.

5 Table 8 shows extracts of the NPTs of the ‘Repair Failure’ variable and the amount of data available
6 to learn its parameters. The values written in bold and italic fonts are the parameters learned by
7 combining the results of the meta-analysis with the data, and the values written in normal fonts are the
8 parameters learned purely from the data. The results of these approaches differ substantially as the
9 data gets smaller. The effects of this difference to the model performance are discussed in Section 5.1.

10 4.4.2 Latent Variables

11 The BN contained several latent variables as described in Section 4.3 (see Table 6 for a list of these
12 variables). Ranked nodes were used to model the NPT of these variables [24]. A ranked node is an
13 approximation of the truncated normal distribution to the multinomial distribution with ordinal scale.
14 We used the framework proposed by Fenton et al. [24] to elicit the parameters of ranked nodes. For
15 each of the latent variables we first asked the domain experts to describe the relation between the
16 variable and its parents. Afterwards, we selected a suitable ranked node function and elicited initial
17 weights that imitate the described relation. We presented the behaviour of the ranked node under
18 various combinations of observations to the domain experts, and refined the weights based on their
19 comments.

20 4.4.3 Variables with Adequate Amount of Data

21 After the parameters with insufficient or no data were defined, the remainder of the parameters were
22 learned purely from the data. The expectation-maximisation (EM) algorithm [25] was used to learn

1 those parameters as the dataset contained missing values. Small correction factors were used at the
2 maximisation step of the EM algorithm to avoid zero probabilities. The parameters that were already
3 defined in the previous steps were kept fixed while EM was applied. For example, EM was applied to
4 the parameters of the ‘Arterial Repair’ variable with more than 20 instances of data. The other
5 parameters of this variable had been already learnt by using meta-analysis results at this stage, so they
6 were kept fixed while the EM was applied.

7 5 Results

8 The performance of the LEVT BN was evaluated using a 10-fold cross-validation [26]. We used
9 multiple performance measures to assess the discrimination, calibration and accuracy of the model.
10 Receiver operating characteristic (ROC) curves, and sensitivity and specificity values was used to
11 assess the discrimination, the Hosmer-Lemeshow (HL) [27] test was used to assess the calibration.
12 The HL test divides the data into multiple subgroups, and calculates a chi-square statistic comparing
13 the observed outcomes to the outcomes expected by the model in each subgroup. Low p-values
14 indicate a lack of calibration. In large datasets, small differences between the expected and observed
15 outcomes can lead to low p-values in a HL test but the visual representation of this test provides a
16 concise summary of the model calibration. The Brier score (BS) and Brier skill score (BSS) was used
17 to assess the accuracy [28,29]. BS is the mean squared difference between the predicted probability
18 and observed outcome. A BS of 0 indicates a perfect model and 1 is the worst score achievable. BSS
19 measures the improvement of the model’s prediction relative to the average probability of the event in
20 the data or another reference probability. A BSS of 1 indicates a perfect model, and a negative value
21 indicates a worse prediction than the average probability. The area under the ROC curve (AUROC) of
22 the LEVT BN was 0.90. When operated at 80% and 90% sensitivity, the specificity was 80% and
23 70% respectively. BS and BSS of the LEVT BN were 0.06 and 0.33 respectively. The BN was well
24 calibrated with a HL statistic of 12.7 (p-value: 0.13). Figure 11 shows a graphical representation of
25 the HL tests. The predictions of the model are divided into deciles in this figure; the expected number
26 of outcomes according to the model and the number of outcomes in the data are compared for each
27 decile.




Figure 11 Calibration of the LEVT BN

5.1 Not Compensating for the Lack of Data


The main reason of using our method for parameter learning was that the available data were insufficient to learn parts of the BN structure. Since the BN model in Figure 8 has a complicated structure compared to the available data, a purely data driven parameter learning approach would probably overfit the model. To show the consequences of not doing anything to compensate for the lack of data, we learned the parameters of the same BN structure from data without using any information from the meta-analysis. The parameters of the variables that had no data (i.e. ranked nodes, see Section 4.4.3) were defined using the same parameters values elicited from experts. The data-driven parameter learning algorithm had poor discrimination, calibration and accuracy. The AUROC was 0.68, the specificity was 29% and 45% at the 90% and 80% sensitivity levels, and the HL test indicated poor calibration (p-value: 0.01). The BS and BSS were 0.10 and 0.02. Our method outperformed the data-driven parameter learning algorithm in all measures. In summary, the purely data driven parameter learning overfitted the training data as the data was inadequate to learn some parts of the BN parameters (Table 7). This underlines the need to exploit other sources of information such as published evidence. Our method overcomes the overfitting problem by using information from the meta-analysis when the data is not enough.

5.2 Mangled Extremity Severity Score

The mangled extremity severity score (MESS) [30] is a well-known scoring system [31] that was developed to provide decision support in the management of patients with severe lower extremity injuries. MESS calculates a score based on the injury mechanism, the degree of shock, the ischaemic status and the patient's age. If the score is above a certain threshold value MESS recommends an

1 early amputation. Our method outperformed MESS in predicting the ‘Nonviable Extremity’ variable
2 (Figure 12). MESS had an AUROC of 0.75. When operated at 90% and 80% sensitivity, its specificity
3 was 40% and 60% respectively. The HL test indicated poor calibration (p-value=0.01). BS and BSS
4 of MESS could not be calculated as its outputs are not probabilities.

5



6

7 *Figure 12 ROC Curves for the LEVT BN and MESS*

8 5.3 Purely Data Driven Structure Learning

9 We compared the performance of our method to three different data-driven structure learning
10 algorithms. These algorithms learn both the structure and parameters from data so they do not use
11 expert knowledge at all in model development. These algorithms avoid overfitting by penalising large
12 BN structures. The following structure algorithms were used:

- 13 1. *A score based learning algorithm:* hill climbing (HC) algorithm with the BIC score [32–34]
- 14 2. *A constraint based algorithm:* grow shrink (GS) algorithm [33]
- 15 3. *A combination of score and constraint based approaches:* max-min hill climbing (MMHC)
16 algorithm [35]
- 17 4. *A score based EM algorithm:* structural EM (SEM) algorithm with the BIC score [36].

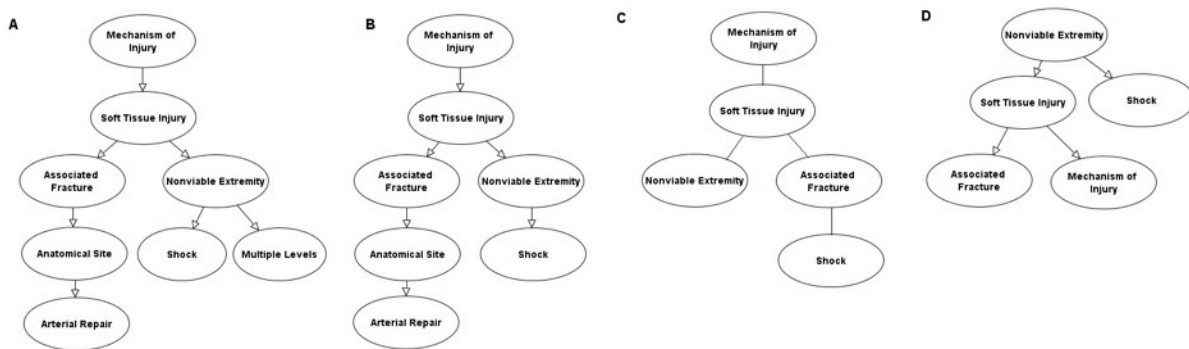
18 The first three algorithms require complete datasets so we imputed the missing values in the dataset
19 using the Amelia package [37] in the R statistical software for these algorithms (see [37] for a

1 description of the imputation technique used in Amelia). The HC, GS and MMHC algorithms are
 2 readily implemented in the BNLearn package of R [38]. The SEM algorithm is able to handle missing
 3 values when learning structure therefore it is not necessary to impute the missing values beforehand.
 4 We used the SEM algorithm implemented in the structure learning package of the Bayes net toolbox
 5 of Matlab [39].

6 *Table 9 Results of our method and the structure learning algorithms*

	LEVT BN	HC	MMHC	GS	SEM
AUROC	0.91	0.83	0.83	0.84	0.84
Specificity (at 90% Sensitivity)	70%	37%	41%	54%	43%
Specificity (at 80% Sensitivity)	81%	71%	80%	69%	81%
Hosmer-Lemeshow Test	12.7 (p=0.13)	11.5 (p=0.17)	15.3 (p=0.05)	8.5 (p=0.39)	13.5 (p=0.10)
Brier Score	0.06	0.07	0.07	0.07	0.07
Brier Skill Score	0.33	0.22	0.21	0.23	0.26

7 The LEVT BN had a better AUROC and substantially better performance at the operating points with
 8 higher sensitivity levels (see Table 9). The LEVT BN had better BS and BSS than the structure
 9 learning algorithms. The LEVT BN and BNs developed by HC and GS approaches were well
 10 calibrated. Since the amount of data was insufficient to learn the relation between some variables, the
 11 structure learning algorithms avoided overfitting by learning simple BN structures. The BN structures
 12 learned by HC, MMHC, GS and SEM algorithms are shown in Figure 13.



13

14 *Figure 13. BN structures learned by (A) HC, (B) MMHC, (C) GS* and (D) SEM algorithms*
 15 **GS algorithm learns the structure with undirected arcs.*

6 Conclusion

This paper presented a novel methodology to build BN decision support models from the results of a meta-analysis, expert knowledge and data. The main contribution of this methodology was a novel parameter learning technique that combines univariate statistics with data to learn multivariate BN parameters. Our method was successfully applied to a trauma case-study of severely injured lower extremities. We developed a BN model that accurately predicts the viability of a lower extremity with vascular trauma. The case study demonstrated the benefits of integrating different sources of evidence into BN development. In a 10-fold cross-validation, the BN built by our approach outperformed the MESS scoring system and four different data-driven structure learning techniques. The AUROC of the LEVT BN was 0.91; whereas it was 0.84 for the best performing structure learning technique, 0.75 for MESS, and 0.68 when meta-analysis values were not used for defining the LEVT BN parameters.

The techniques presented in this paper can be applied to a wider scope of problems than trauma care. Using our method, models that reflect complexity of clinical decisions can be built even when there is insufficient patient data. Our method enables the use of information from other sources such as published evidence and meta-analysis of systematic reviews. It offers a Bayesian way of combining multivariate patient data with published statistics conditioned on a smaller number of variables.

As further research, our auxiliary parameter learning model could be expanded to a unified parameter learning framework that combines data, published evidence and domain knowledge. Qualitative expert constraints [7,40,41] could be integrated into our parameter learning method to incorporate expert knowledge alongside data and meta-analysis. Moreover, the variance estimated from the auxiliary parameter learning model (Section 3.2) could be used to show how well the parameters are understood. In our case study, we applied our parameter learning technique to the parameters that had less than 20 instances of relevant data. The effects using different data thresholds could be explored. Finally, Bayesian parameter learning methods estimate the entire probability distribution of a parameter. The expected value of this distribution is used for the relevant NPT but the variance is

1 often ignored. Ways of integrating the variance to inference and parameter estimation techniques
2 could be investigated.

3 Acknowledgments

4 We thank two anonymous reviewers for their valuable comments and suggestions that have helped us
5 to improve this paper.

6 References

- 7 1. Buchan I, Winn J, Bishop C (2009) A unified modelling approach to data intensive healthcare.
8 The fourth paradigm: data-intensive scientific discovery. Redmond, WA: Microsoft Research.
9 pp. 91–97.
- 10 2. Marshall JC (2006) Surgical Decision-Making: Integrating Evidence, Inference, and Experience.
11 Surgical Clinics of North America 86: 201–215. doi:10.1016/j.suc.2005.10.009.
- 12 3. Moons KGM, Royston P, Vergouwe Y, Grobbee DE, Altman DG (2009) Prognosis and
13 prognostic research: what, why, and how? BMJ 338: b375–b375. doi:10.1136/bmj.b375.
- 14 4. Daly R, Shen Q, Aitken S (2011) Learning Bayesian networks: approaches and issues.
15 Knowledge Engineering Review 26: 99–157.
- 16 5. Flores JM, Nicholson AE, Brunskill A, Korb KB, Mascaro S (2011) Incorporating expert
17 knowledge when learning Bayesian network structure: A medical case study. Artificial
18 intelligence in medicine 53: 181–204.
- 19 6. Yet B, Perkins Z, Fenton N, Tai N, Marsh W (2014) Not just data: A method for improving
20 prediction with knowledge. Journal of biomedical informatics 48: 28–37.
21 doi:10.1016/j.jbi.2013.10.012.
- 22 7. Zhou Y, Fenton N, Neil M (2014) Bayesian network approach to multinomial parameter
23 learning using data and expert judgments. International Journal of Approximate Reasoning.
24 doi:10.1016/j.ijar.2014.02.008.
- 25 8. Van Houwelingen HC, Arends LR, Stijnen T (2002) Advanced methods in meta-analysis:
26 multivariate approach and meta-regression. Statistics in medicine 21: 589–624.
- 27 9. Vickers AJ (2006) Whose data set is it anyway? Sharing raw data from randomized. Trials 7:
28 15. doi:10.1186/1745-6215-7-15.
- 29 10. Rawlins M (2008) De testimonio: on the evidence for decisions about the use of therapeutic
30 interventions. Clinical Medicine 8: 579–588.
- 31 11. Greenhalgh T, Howick J, Maskrey N, for the Evidence Based Medicine Renaissance Group
32 (2014) Evidence based medicine: a movement in crisis? BMJ 348: g3725–g3725.
33 doi:10.1136/bmj.g3725.
- 34 12. Druzdzel MJ, Díez FJ (2003) Combining knowledge from different sources in causal
35 probabilistic models. The Journal of Machine Learning Research 4: 295–316.

- 1 13. Lappenschaar M, Hommersom A, Lucas PJF, Lagro J, Visscher S (2013) Multilevel Bayesian
2 networks for the analysis of hierarchical health care data. *Artificial Intelligence in Medicine* 57:
3 171–183. doi:10.1016/j.artmed.2012.12.007.
- 4 14. Higgins J, Thompson SG, Spiegelhalter DJ (2009) A re-evaluation of random-effects meta-
5 analysis. *Journal of the Royal Statistical Society: Series A (Statistics in Society)* 172: 137–159.
- 6 15. Müller M, Wandel S, Colebunders R, Attia S, Furrer H, et al. (2010) Immune reconstitution
7 inflammatory syndrome in patients starting antiretroviral therapy for HIV infection: a systematic
8 review and meta-analysis. *The Lancet Infectious Diseases* 10: 251–261. doi:10.1016/S1473-
9 3099(10)70026-8.
- 10 16. Spiegelhalter DJ, Abrams KR, Myles JP (2004) Bayesian approaches to clinical trials and
11 health-care evaluation. John Wiley & Sons.
- 12 17. Neil M, Tailor M, Marquez D (2007) Inference in hybrid Bayesian networks using dynamic
13 discretization. *Statistics and Computing* 17: 219–233.
- 14 18. Agena Ltd (2014) AgenaRisk: Bayesian Network and Simulation Software for Risk Analysis
15 and Decision Support. Available: <http://www.agenarisk.com>. Accessed 15 June 2014.
- 16 19. Lunn D, Spiegelhalter D, Thomas A, Best N (2009) The BUGS project: Evolution, critique and
17 future directions. *Statistics in medicine* 28: 3049–3067.
- 18 20. Fenton NE, Neil MD (2012) Risk Assessment and Decision Analysis with Bayesian Networks.
19 CRC Press.
- 20 21. Koller D, Friedman N (2009) Probabilistic graphical models principles and techniques.
21 Cambridge (Mass.): MIT Press.
- 22 22. Perkins Z, Glasgow, Tai N (2012) A systematic review of prognostic factors related to
23 secondary amputation in patients with lower limb vascular trauma requiring surgical repair.
24 PROSPERO. Available:
25 http://www.crd.york.ac.uk/PROSPERO/display_record.asp?ID=CRD42012002720. Accessed
26 14 March 2014.
- 27 23. Gelman A, Rubin DB (1992) Inference from iterative simulation using multiple sequences.
28 *Statistical science*: 457–472.
- 29 24. Fenton NE, Neil M, Caballero JG (2007) Using ranked nodes to model qualitative judgments in
30 Bayesian networks. *Knowledge and Data Engineering, IEEE Transactions on* 19: 1420–1432.
- 31 25. Lauritzen SL (1995) The EM algorithm for graphical association models with missing data.
32 *Computational Statistics & Data Analysis* 19: 191–201.
- 33 26. Kohavi R (1995) A Study of Cross-validation and Bootstrap for Accuracy Estimation and
34 Model Selection. *Proceedings of the 14th International Joint Conference on Artificial*
35 *Intelligence - Volume 2. IJCAI'95. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc.*
36 pp. 1137–1143.
- 37 27. Hosmer DW, Lemeshow S (1980) Goodness of fit tests for the multiple logistic regression
38 model. *Communications in statistics Theory and methods* 9: 1043–1069.
- 39 28. Brier GW (1950) Verification of forecasts expressed in terms of probability. *Monthly weather*
40 *review* 78: 1–3.

- 1 29. Weigel AP, Liniger MA, Appenzeller C (2007) The discrete Brier and ranked probability skill
2 scores. *Monthly Weather Review* 135: 118–124.
- 3 30. Johansen KAJ, Daines M, Howey T, Helfet D, Hansen Jr ST (1990) Objective criteria
4 accurately predict amputation following lower extremity trauma. *The Journal of Trauma and*
5 *Acute Care Surgery* 30: 568–573.
- 6 31. Bosse MJ, MacKenzie EJ, Kellam JF, Burgess AR, Webb LX, et al. (2001) A prospective
7 evaluation of the clinical utility of the lower-extremity injury-severity scores. *The Journal of*
8 *Bone & Joint Surgery* 83: 3–14.
- 9 32. Korb KB, Nicholson AE (2004) *Bayesian artificial intelligence*. CRC Press.
- 10 33. Margaritis D (2003) *Learning Bayesian network model structure from data* University of
11 Pittsburgh. Available: <http://www.cs.cmu.edu/afs/cs.cmu.edu/user/dmarg/www/Papers/PhD->
12 [Thesis-Margaritis.pdf](http://www.cs.cmu.edu/afs/cs.cmu.edu/user/dmarg/www/Papers/PhD-). Accessed 14 March 2014.
- 13 34. Schwarz G (1978) Estimating the dimension of a model. *The annals of statistics* 6: 461–464.
- 14 35. Tsamardinos I, Brown LE, Aliferis CF (2006) The max-min hill-climbing Bayesian network
15 structure learning algorithm. *Machine learning* 65: 31–78.
- 16 36. Friedman N (1998) The Bayesian structural EM algorithm. *Proceedings of the Fourteenth*
17 *conference on Uncertainty in artificial intelligence*. pp. 129–138. Available:
18 <http://dl.acm.org/citation.cfm?id=2074110>. Accessed 28 May 2013.
- 19 37. Honaker J, Gary King, Blackwell M (2014) *Amelia II: A Program for Missing Data*. Available:
20 <http://cran.r-project.org/web/packages/Amelia/vignettes/amelia.pdf>. Accessed 15 July 2014.
- 21 38. Scutari M (2010) *Learning Bayesian networks with the bnlearn R package*. *Journal of Statistical*
22 *Software* 35: 1–29.
- 23 39. Francois O (2010) *Structure Learning Package for BayesNet Toolbox*. Available:
24 <http://ofrancois.tuxfamily.org/slp.html>. Accessed 15 July 2014.
- 25 40. Feelders A, Van der Gaag LC (2006) Learning Bayesian network parameters under order
26 constraints. *International Journal of Approximate Reasoning* 42: 37–53.
- 27 41. Tong Y, Ji Q (2008) Learning Bayesian networks with qualitative constraints. *IEEE Conference*
28 *on Computer Vision and Pattern Recognition, 2008. (CVPR 2008)*. pp. 1–8.
29 doi:10.1109/CVPR.2008.4587368.

30