

Embellishing a Bayesian Network using a Chain Event Graph

L. M. Barclay*, J. L. Hutton, J. Q. Smith

Department of Statistics, University of Warwick, Coventry CV4 7AL, United Kingdom

Abstract

The search for a useful explanatory model based on a Bayesian Network (BN) now has a long and successful history. However, when the dependence structure between the variables of the problem is asymmetric then this cannot be captured by the BN. The Chain Event Graph (CEG) provides a richer class of models which incorporates these types of dependence structures as well as retaining the property that conclusions can be easily read back to the client. We demonstrate on a real health study how the CEG leads us to promising higher scoring models and further enables us to make more refined conclusions than can be made from the BN. Further we show how these graphs can express causal hypotheses about possible intervention that could be enforced.

Keywords: Chain Event Graphs, Bayesian Networks, Bayesian Model Selection, Causality

1. Introduction

The Bayesian Network (BN) is the most widely used graphical model (Cowell et al., 2007; Korb and Nicholson, 2004) which expresses the relationship between the variables of the system in terms of conditional independence statements. Its graphical structure makes it a particularly useful tool to feed conclusions back to the client and it has therefore been employed in many real-world applications. However, in certain cases the BN does not provide a rich enough structure to incorporate all information obtainable from the data set. This is the case, for example, when the conditional independence statements of the problem are asymmetric or only certain combinations of variables affect another variable and this cannot be represented simply by the directed edges between variables in the BN (Poole and Zhang, 2003). To take these features into account extensions to the BN have been proposed, mostly in the form of tables or tree-like structures which are added to the graph, leading to the context-specific Bayesian Network (Boutilier et al., 1996; Poole and Zhang, 2003; Friedman and Goldszmidt, 1998). However, these methods focus primarily on efficient propagation and learning and lose the benefit of the BN's expressiveness for the client.

*Corresponding author

Email addresses: L.M.Barclay@warwick.ac.uk (L. M. Barclay), J.L.Hutton@warwick.ac.uk (J. L. Hutton), J.Q.Smith@warwick.ac.uk (J. Q. Smith)

The Chain Event Graph (CEG) is a new flexible class of graphical models which can represent asymmetric structures directly in its topology. It is related to the probability decision graph (Oliver, 1993; Jaeger, 2004) and retains the framework of a probability tree in a more compact graph. Because of its graphical derivation it inherits many of the benefits of a BN. For example, we can read off conditional independence statements directly from the topology of the graph (Smith and Anderson, 2008), carry out model selection on CEGs (Freeman and Smith, 2011) and run fast propagation algorithms (Thwaites et al., 2008). It also admits a causal extension (Thwaites et al., 2010; Thwaites, 2012). On the other hand it is a more general class than either the probability decision graph or the discrete BN containing each of these classes as a special case (Smith and Anderson, 2008).

In this paper we illustrate how we use the CEG to embellish an initial BN model description, demonstrating how this can provide a more detailed but still transparent explanation of an underlying process, which seamlessly enhances an original BN analysis. The application in this paper explores a particular study of the Christchurch Health and Development Study (web, 2012), looking at the effect of social and family factors on children's health in a New Zealand birth cohort (Fergusson et al., 1986).

We begin the paper with a description of the programme of study forming the basis of this analysis. In section 3 we report how an initial routine BN model search enabled us to elaborate the conclusions made in the original study. In sections 4 and 5 we show how we then made more detailed exploration using the CEG, discovering a much better fitting model that - like a BN - could be read back to the client in narrative style. We also illustrate how the fitted model can be given a causal interpretation which, if valid, allows us to perform various 'what - if' analyses under various policy controls. In the final section we propose further extensions to the present analysis.

2. The Christchurch Health and Development Study

The Christchurch Health and Development Study (CHDS) is carried out by a research group at Otago University, led by Professor David Fergusson. It is a cohort study which has followed up 1265 children born in mid 1977 in Christchurch, New Zealand, for over 30 years. The CHDS has explored the children's development from childhood to adulthood regarding their education, behaviour and health with respect to a wide range of social, economic and family factors.

The running example used in this paper reanalyses an early subset of the CHDS discussed in Fergusson et al. (1986) which studies the first five years of the Christchurch cohort looking at the effect the family's social background, the economic status and the number of family life events have on the child's health which is measured by rates of hospital admission.

Based on previous analysis, Fergusson et al. concluded to only consider admissions due to illness and accidents as these were the only reasons for admissions that are sensitive to social and family situation. To describe the family's social background the CHDS group collected information about the mother's education and age at birth, the family's socioeconomic status and ethnic origin, and whether the child grew up in a single or two parent family. These variables were then combined using factor analysis to give a single measure of the social background (see Fergusson et al. (1984) for details). Similarly, the economic status was measured as a function

of the family income, possible financial difficulties and by rating the standard of living and the quality of the accommodation of the child. Again these were simplified into a single measure of the overall economic situation. Of particular interest here was whether the effect of adverse life events in the children's lives might be associated with an increased rate of illness. This has provoked controversy and continues to be a subject of research (Weich et al., 2009). Twenty events were classed as life events. These included the experience of moving house, a husband changing job, the death of a close friend or relative, serious financial problems, divorce or a serious illness or accident (Beautrais et al., 1982).

Data was collected through a variety of sources: Mothers were interviewed every year on the family life events, while the economic situation was predominantly rated every year by an interviewer. Hospital admissions were also obtained by interviewing the mother of every child at each year and further by asking them to keep a health diary record. The results were additionally compared with the central records of the Christchurch Hospital.

Rates of hospital admissions were first compared for the various levels of each covariate separately using one-way analysis of variance, concluding that the more socially or economically disadvantaged the child's background (without adjusting for the other factors) the higher the hospital admission rate tended to be. In addition to this the probability of an admission increased significantly with the number of family life events. For further analysis a Cox proportional hazards model estimating the risk of hospital admission over the five years was fitted. Results from Fergusson et al. (1986) showed that according to this model the family's economic status did not influence the risk of admission significantly once adjusting for the other covariates, suggesting that in this type of population financial problems were not the main reason for health problems. In contrast, family life events and social background both appeared to have a significant impact on the admission rates, even once having adjusted for the other covariates. In particular, the most significant association was found between the hospital admissions rate and the number of family life events. These interesting conclusions motivated the following analysis. We demonstrate that our graphical methods not only allow us to disentangle the association between the variables but enable us to present our results in a transparent way.

We aimed to follow as far as possible the original study outlined above: Complete data was available to us for 890 children and so our analysis was carried out on these. Because the variables describing the social and economic background are discrete, predominantly with few categories, we slightly adapted the methods of Fergusson et al. (1986) who used a factor analysis. Instead we fitted a latent-class model to construct a categorical variable for the social background and the economic situation. The predicted classes of each child were used to describe its social and economic situation. For simplicity, we here assumed binary latent classes throughout, distinguishing between a high or low social background and a high or low economic situation. We further made a distinction between 'no hospital admission' and 'at least one hospital admission' since the counts for more than one admission are sparse. Similarly, we divided the life events into the three approximately equal categories: 'low' (0 – 5 events over the five years), 'average' (6 – 9 events) and 'high' (at least 10 events). There were 169 (19%) children with at least one admission and the proportion of admission ranged from 0.12 to 0.26 by background (Table 1).

Admissions	No Admission 721	≥ 1 Admission 169	
Social Background	High 507 (0.148)	Low 383 (0.245)	
Economic Situation	High 283 (0.148)	Low 607 (0.209)	
Number of Life Events	0 – 5 Events 329 (0.119)	6 – 9 Events 295 (0.210)	≥ 10 Events 266 (0.256)

Table 1: Summary statistics and proportions of admission

3. Translation into a Bayesian Network versus model selection

We began by fitting a standard BN to this application outlined in the previous section.

We let

- X_1 = family social position
- X_2 = family economic situation
- X_3 = number of family life events
- X_4 = hospital admission.

The main conclusions of Fergusson et al. (1986) could then be expressed as the single conditional independence statement $X_4 \perp\!\!\!\perp X_2 \mid X_1, X_3$ or equivalently by the DAG given in Figure 1.




Figure 1: DAG for the effects on hospital admission

Our first step was to carry out a standard model selection procedure to find the Maximum a Posteriori (MAP) BN structure over the four variables. The network structures were scored using the Bayesian Dirichlet (BD) metric in Heckerman et al. (1995). As recommended by Neapolitan (2004) the prior Dirichlet distributions were given an equivalent sample size of 3, the highest number of categories taken by a variable in the problem. For transparency, we chose a prior network such that the distribution over all possible configurations was uniform and assumed that structures were a priori equally likely. We use Bayes Factor (BF) scores relative to the MAP BN model throughout for the comparison of different models although other scoring methods could have been used.

An exhaustive search using the ‘deal’ package in R (Bøttcher and Dethlefsen, 2003) over all possible BNs on the four variables found the MAP model to be the DAG given in Figure 2(a)

with the corresponding set of conditional independence statements given by: $X_3 \perp\!\!\!\perp X_2 \mid X_1$, $X_4 \perp\!\!\!\perp X_1, X_2 \mid X_3$ and the associated predictive conditional probability vectors given in Table 2.

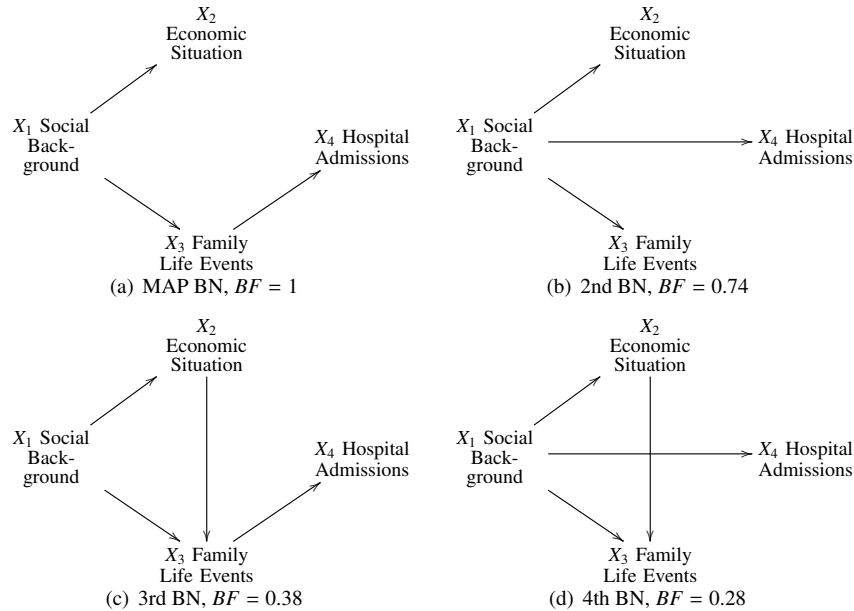


Figure 2: Highest scoring BN structures

Conditional Probability Vector	
$(P(X_1 = \text{High}), P(X_1 = \text{Low}))$	(0.569, 0.431)
$(P(X_2 = \text{High} \mid X_1 = \text{High}), P(X_2 = \text{Low} \mid X_1 = \text{High}))$	(0.468, 0.532)
$(P(X_2 = \text{High} \mid X_1 = \text{Low}), P(X_2 = \text{Low} \mid X_1 = \text{Low}))$	(0.122, 0.878)
$(P(X_3 = \text{Low} \mid X_1 = \text{High}), P(X_3 = \text{Average} \mid X_1 = \text{High}), P(X_3 = \text{High} \mid X_1 = \text{High}))$	(0.461, 0.347, 0.192)
$(P(X_3 = \text{Low} \mid X_1 = \text{Low}), P(X_3 = \text{Average} \mid X_1 = \text{Low}), P(X_3 = \text{High} \mid X_1 = \text{Low}))$	(0.248, 0.311, 0.441)
$(P(X_4 = \text{No admission} \mid X_3 = \text{Low}), P(X_4 = \text{Admission} \mid X_3 = \text{Low}))$	(0.881, 0.119)
$(P(X_4 = \text{No admission} \mid X_3 = \text{Average}), P(X_4 = \text{Admission} \mid X_3 = \text{Average}))$	(0.789, 0.211)
$(P(X_4 = \text{No admission} \mid X_3 = \text{High}), P(X_4 = \text{Admission} \mid X_3 = \text{High}))$	(0.744, 0.256)

Table 2: Conditional Probability Table for MAP BN

In comparison with the network structure derived from Fergusson et al. (1986) the MAP model found was more sparse: It suggests an additional conditional independency between the economic situation and the family life events given the social background ($X_3 \perp\!\!\!\perp X_2 \mid X_1$) and further expresses a direct dependency only between the life events and the hospital admissions and not between social background and admissions ($X_4 \perp\!\!\!\perp X_1, X_2 \mid X_3$). Further, the exhaustive search over all possible structures revealed three further network structures scoring only slightly less than the MAP model, given in Figure 2 (b)-(d). Network structure (b) swaps the directed edge from family life events to admissions with an edge from the social background to the admissions. Structure (c) introduces an extra edge between the economic situation and the family life events and model (d) again exchanges the edge of model (c) from family life events to admissions for an edge between social background and admissions. This discrepancy in the dependence struc-

ture of the variables lets us believe that there may be a more complex relationship between the variables that cannot be captured by the BN.

In such applications we are interested in the causal effect of social background, economic situation and family life events on hospital admissions. In this paper we define a cause in the sense of Pearl (2009) in terms of interventions, where a variable X_1 is causal for a variable X_2 if manipulation with respect to X_1 affects the distribution of X_2 . If we are prepared to treat the Bayesian Networks above causally, then Pearl's intervention formula (Pearl, 2009) enables us to calculate the predicted effect of an intervention on the hospital admissions. Thus if it were possible to enact a policy ensuring that the life events on a particular unit would lie in the 'low' category, then the causal hypothesis above would allow us to conclude that the consequent probability of hospital admission would be 0.119 (compare Table 2). However, these types of interventions may also be asymmetric. We may be interested in the effect of an intervention which gives only families from a low social background financial aid, or similarly, we would like to know the predicted probability of an admission if we could stop children from a low social background having a high number of life events.

4. Embellishing a Bayesian Network using a Chain Event Graph

We therefore introduce a new, more flexible, class of models called the Chain Event Graphs (CEG) (Smith and Anderson, 2008; Thwaites et al., 2010). It has several advantages over the BN: Firstly, it is an embellishment of the standard BN by allowing for asymmetries within the dependence structure as well as capturing discrete BNs as a particular subclass. Secondly, the CEG is derived from a probability tree retaining its paths in a more compact graph. It hence enables us to provide a plausible story of the way in which different factors affect children's health, sharing therefore with the BN the property of providing an evocative graphical framework through which conclusions can be read back to the client.

The CEG simplifies the tree by introducing the concepts of 'stages' and 'positions' which group the vertices in the tree together. We say that two vertices are in the same stage u , when we have a one-to-one mapping between the edges emanating from the two vertices and their associated conditional probabilities are the same. The full set of stages is denoted by $J(T)$. When two vertices are in the same stage then their associated edges are coloured such that corresponding edges have the same colour. The resulting tree is called a staged tree. For clarity, we here use instead of colours different node shapes when two vertices are in the same stage. A finer partition of the vertices is given by the concept of positions. We say that two vertices are in the same position w when their full subtrees and all associated probabilities are the same.

The CEG is then given by the staged tree collapsed over the positions of the tree, with the set of leaf-nodes collected in a single position, called w_∞ . Hence it is a graphical model whose vertices equal the positions of the corresponding staged tree, and its edge set is defined as follows: For each position w a single representative vertex $v(w)$ is chosen. Then there exists an edge from a position w to another position w' in the CEG for every edge in the staged tree from $v(w)$ to a vertex $v' \in w'$. As a stage may be made up of several positions, the stages of the tree are indicated by adding an undirected dotted line between any two positions that are in the same stage.

We illustrate below how a discrete BN can be represented as a CEG using the MAP BN structure given in Figure 2(a). For a detailed proof that every discrete BN can be written as a CEG see

Smith and Anderson (2008). We can first draw a tree corresponding to our discrete BN such that parent variables appear before their children in the ordering of the tree. Note that this ordering is not necessarily unique. The conditional independence statements of the BN can then be uniquely represented by defining stages on this tree. For our example we have the corresponding staged tree given in Figure 3 with stages:

$$u_0 = \{v_0\}, u_1 = \{v_1\}, u_2 = \{v_2\}, u_3 = \{v_3, v_4\}, u_4 = \{v_5, v_6\}, \\ u_5 = \{v_7, v_{10}, v_{13}, v_{16}\}, u_6 = \{v_8, v_{11}, v_{14}, v_{17}\}, u_7 = \{v_9, v_{12}, v_{15}, v_{18}\}.$$

We have that no conditional independencies in the BN corresponds to a CEG where each node is in a separate stage. A conditional independency in the BN of the form

$$X_i \perp\!\!\!\perp \{X_1, \dots, X_{i-1}\} \setminus pa(X_i) \mid pa(X_i),$$

where $pa(X_i)$ is the parent set of X_i , then puts all the vertices into the same stage whose previous events only differ in $\{X_1, \dots, X_{i-1}\} \setminus pa(X_i)$. We hence have a one-to-one correspondence between the stages of the tree and the parent configurations $pa(x_i)$, $i = 2, 3, 4$, of the BN and so every BN can be written as a staged tree.

From the staged tree we can then simply construct the CEG. In our example the stages and positions coincide, and hence the vertices of the CEG are

$$w_0 = \{v_0\}, w_1 = \{v_1\}, w_2 = \{v_2\}, w_3 = \{v_3, v_4\}, w_4 = \{v_5, v_6\}, \\ w_5 = \{v_7, v_{10}, v_{13}, v_{16}\}, w_6 = \{v_8, v_{11}, v_{14}, v_{17}\}, w_7 = \{v_9, v_{12}, v_{15}, v_{18}\}, \\ w_\infty = \{v_{19}, \dots, v_{42}\},$$

with the corresponding graph given in Figure 4.

It can be seen from Figure 4 that the CEG of a BN is very symmetric. The increased generality of the CEG, however, allows us to search a much larger space of graphically explanatory models to find higher scoring models with less symmetry. We note that the loss of symmetry does not make the resulting fitted model any less interpretable, as will be illustrated in the next section.

5. Scoring Bayesian Networks and Chain Event Graphs

When searching the CEG space we can find the most probable BN derived previously, as the BN is a subclass of CEGs. However, if we consider the BN structures in Figure 2 as well as the deductions we made by Fergusson et al. (1986), it seems likely that the CEG will combine certain vertices into stages and positions in an asymmetric way resulting in a better model score. For example, it may be that a child from a lower social background and a low number of life events has the same probability of admissions as a child from a family with a higher social background and an average number of life events.

We set up an equivalent method for scoring all possible CEGs of the problem to allow for a direct comparison between BNs and CEGs. Freeman and Smith (2011) show that the axioms in Heckerman et al. (1995) required for the BD-metric and the set up of the hyperparameters

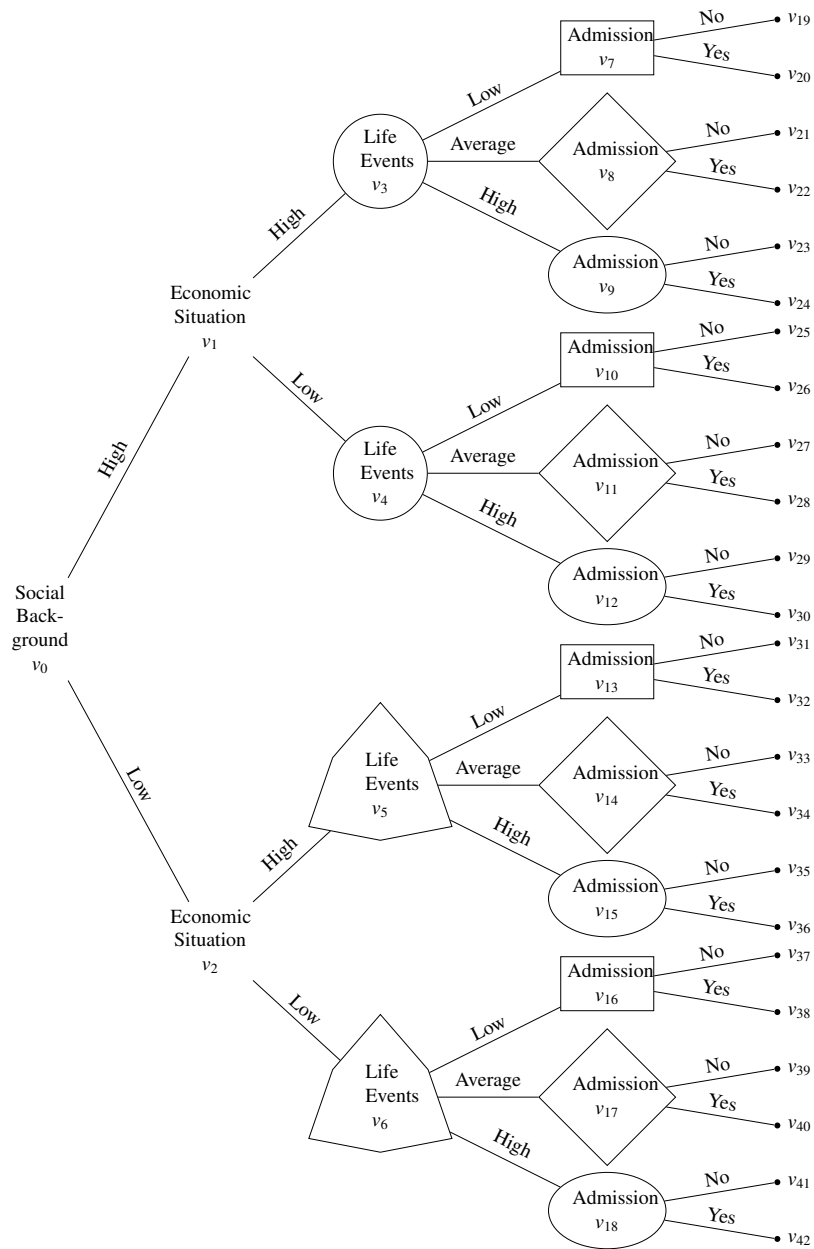


Figure 3: Staged Tree

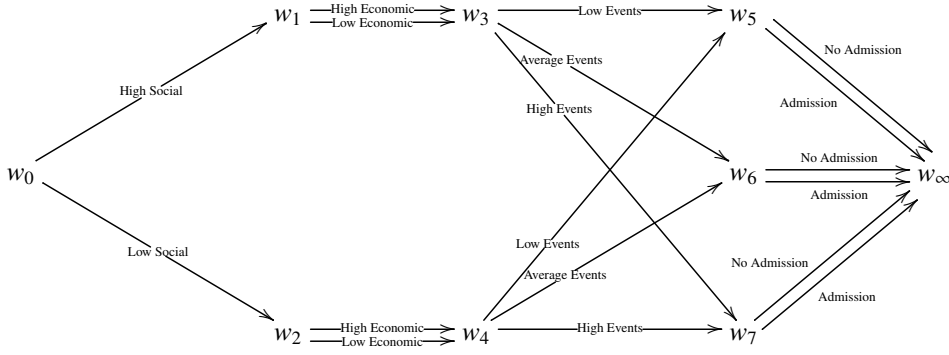


Figure 4: CEG derived from MAP BN structure

of the priors in the BN can be extended in a natural way to the CEG. They prove that under the assumptions that the stage priors are a priori independent and are identical for equivalent stages in different CEG structures, Dirichlet priors on the conditional probability vectors of the staged tree are inevitable. Further priors across models can then be deduced by summing the hyperparameters of the Dirichlet distribution when two stages are merged.

The BD-metric for CEGs is then the posterior probability of a CEG structure C given a complete multinomial sample D which is given explicitly as

$$p(C|D) = \frac{p(C)}{p(D)} p(D|C) = \frac{p(C)}{p(D)} \prod_{u \in J(T)} \frac{\Gamma(\alpha_u)}{\Gamma(\alpha_u + N_u)} \prod_{k=1}^{r_u} \frac{\Gamma(\alpha_{uk} + N_{uk})}{\Gamma(\alpha_{uk})}, \quad (1)$$

with $\alpha_u = \sum_k \alpha_{uk}$ and $N_u = \sum_k N_{uk}$, where α_{uk} are the parameters of the Dirichlet distribution describing the probability of going from a vertex in stage u to a vertex k and, similarly, N_{uk} , $k = 1, \dots, r_u$, is the number of times we observe an individual going from a vertex in stage u to a vertex k .

As for the BN, a uniform prior was given to the root-to-leaf paths of the finest partition of the CEG, in which all non-leaf vertices are in a separate stage. To enable a direct comparison the equivalent sample size was also chosen to be 3, as for the BNs. Assuming that different CEG structures are a priori equally likely, we can hence look at the Bayes Factor scores between BNs and CEGs to compare the model fit. We note that some care is needed when the priors are not chosen to be equally likely as the class of CEGs is larger than the class of BNs.

Since the CEG model space is far larger than the space of BN structures it is not feasible to perform an exhaustive search in all but the simplest case. In our example we therefore implemented a Bayesian Agglomerate Clustering (AHC) Algorithm developed for CEGs in Freeman and Smith (2011). The algorithm starts at the finest partition of the CEG, where each non-leaf vertex is in a separate stage. It then efficiently searches over the model space by finding at every step the two stages, which, when merged, provide the highest CEG score, given by the BD-metric for CEGs. The new model's score and the corresponding CEG structure are recorded at every step. The algorithm stops once the coarsest partition of the CEG has been reached, where all non-leaf vertices with the same topology are in a single stage. The CEG with the highest overall score is then selected.

The AHC algorithm found the highest scoring CEG structure to be the graph in Figure 5. The corresponding stages are:

$$u_0 = \{v_0\}, u_1 = \{v_1\}, u_2 = \{v_2\}, u_3 = \{v_3, v_4, v_5\}, u_5 = \{v_6\},$$

$$u_6 = \{v_7, v_{10}\}, u_8 = \{v_8, v_{11}, v_{13}, v_{14}, v_{16}\}, u_9 = \{v_9, v_{12}, v_{15}, v_{17}, v_{18}\},$$

where u_3 is split into two positions w_3 and w_4 in Figure 5.

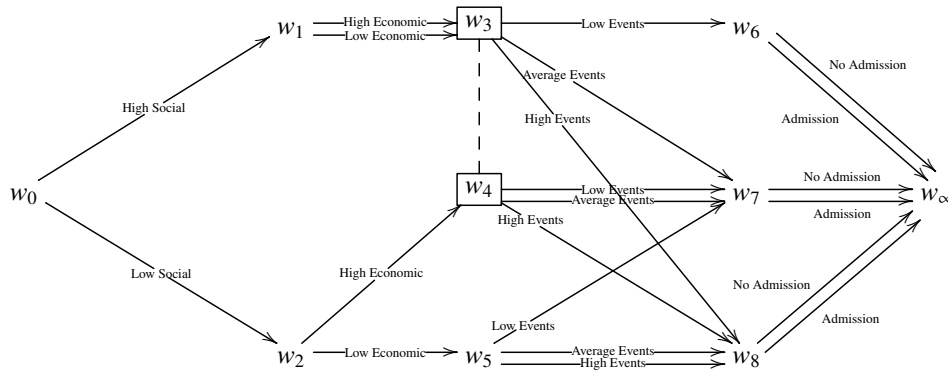


Figure 5: CEG found by the AHC algorithm

The graph illustrates several asymmetries which allow us to draw a number of conclusions on the combined effect of the three covariates on the hospital admissions that are not obtainable from the BN. Below we give a detailed account of the dependence structure of the variables which can be read off directly from the topology of the CEG.

1. The social background appears to have an effect on the economic situation as suggested by the BN ($w_0 \rightarrow w_1, w_0 \rightarrow w_2$).
2. The economic situation seems to have no effect on the number of life events for families from a higher social background, as we have merged situations v_3 and v_4 into w_3 ($w_1 \rightarrow w_3$). However, in a family from a lower social background the economic situation seems to affect the number of life events that occur ($w_2 \rightarrow w_4, w_2 \rightarrow w_5$).
3. Children from a family of high social background and a low number of life events, independent of the economic situation, are in a separate position and hence have a different probability of admissions to the other individuals ($w_3 \rightarrow w_6$).
4. Children from socially advantaged families with an average number of life events are in the same position as children from socially disadvantaged families with a high economic situation and a low or average number of life events, as are children from a low economic situation with a low number of life events ($w_3 \rightarrow w_7, w_4 \rightarrow w_7, w_5 \rightarrow w_7$).
5. All individuals with a high number of life events are in the same position irrespective of their social or economic background. Further, an individual from a low social and economic background with only an average number of life events is also in this position ($w_3 \rightarrow w_8, w_4 \rightarrow w_8, w_5 \rightarrow w_8$).

When comparing this CEG structure with the MAP BN structure, we obtain a Bayes Factor of 79 698 in favour of the CEG, providing very strong evidence that the additional flexibility of

Stage	Conditional Probability Vector
u_0	$(P(X_1 = \text{High}), P(X_1 = \text{Low}))$ (0.569, 0.431)
u_1	$(X_2 = \text{High} u_1) + P(X_2 = \text{Low} u_1)$ 1
u_2	$(P(X_2 = \text{High} u_2), P(X_2 = \text{Low} u_2))$ (0.122, 0.878)
u_3	$(P(X_3 = \text{Low} u_3), P(X_3 = \text{Average} u_3), P(X_3 = \text{High} u_3))$ (0.461, 0.343, 0.196)
u_5	$(P(X_3 = \text{Low} u_5), P(X_3 = \text{Average} u_5), P(X_3 = \text{High} u_5))$ (0.220, 0.312, 0.468)
u_6	$(P(X_4 = \text{No admission} u_6), P(X_4 = \text{Admission} u_6))$ (0.910, 0.090)
u_7	$(P(X_4 = \text{No admission} u_7), P(X_4 = \text{Admission} u_7))$ (0.794, 0.206)
u_8	$(P(X_4 = \text{No admission} u_8), P(X_4 = \text{Admission} u_8))$ (0.743, 0.257)

Table 3: Conditional Probability Table for CEG

the CEG enables us to find a strongly preferable model. Table 3 gives the vectors of predictive probabilities associated with each stage.

Of particular interest are the three final positions w_6 , w_7 and w_8 as these give an interpretation of the effect of a combination of variables on the hospital admissions as described in points 3 to 5 above. The CEG lets us trace the different paths the individuals can take ending up in one of these three positions. We have that the predicted probability of a hospital admission for individuals that reach position w_6 is 0.09. Further we predict an admission probability of 0.206 for individuals that reach position w_7 and 0.257 for the individuals at position w_8 . We note that while a high number of life events forces the individuals into position w_8 with the highest admission probability, an individual from a low social background will never reach position w_6 even when he only has a low number of life events and his economic situation is high. The table of predictive probabilities further illustrates that a child from a low social background is more likely to also have a low economic background. Similarly, a child from a low social and economic background has a predictive probability of 0.461 to have a high number of life events and 0.196 for a low number of life events, while for the remaining children these probabilities are 0.220 and 0.468, respectively. It is also interesting to compare these probabilities with the probability vectors of the MAP BN in Table 2. Our results support the conclusions of Fergusson et al. (1986), that the effect of life events on admissions is strongest. However, the CEG further explains explicitly the way in which the social background and the economic situation may have an additional effect on hospital admissions.

We further illustrate that just as with BNs these graphs can be linked to causal hypotheses about the likely effect of interventions which can then be tested in the future. In particular, the CEG lets us make more refined causal statements about the hospital admissions by allowing for possible interventions to be asymmetric in the sense that we can intervene at a particular position rather than only on a variable. The simplest form of intervention forces an individual that reaches a position w along a particular edge and hence the conditional probability vector at w has one entry equal to 1 and zeros everywhere else. Thwaites et al. (2010) then define, analogous to Pearl (2009), a Causal CEG to be a CEG which, under the manipulation at a position w , obeys the intervention formula that $P(e(w, w^*)|w) = 1$ and all other probabilities are as in the unmanipulated tree. The manipulated CEG can be drawn by deleting all paths emanating from w that do not go along the required edge and we can calculate the effect of the intervention directly from the table of conditional probability vectors.

In our example, if we are prepared to read the graph causally, the effect of giving only families from a low social background financial aid corresponds to forcing all individuals that reach w_2

along the edge $w_2 \rightarrow w_4$. The manipulated CEG is given in Figure 6.

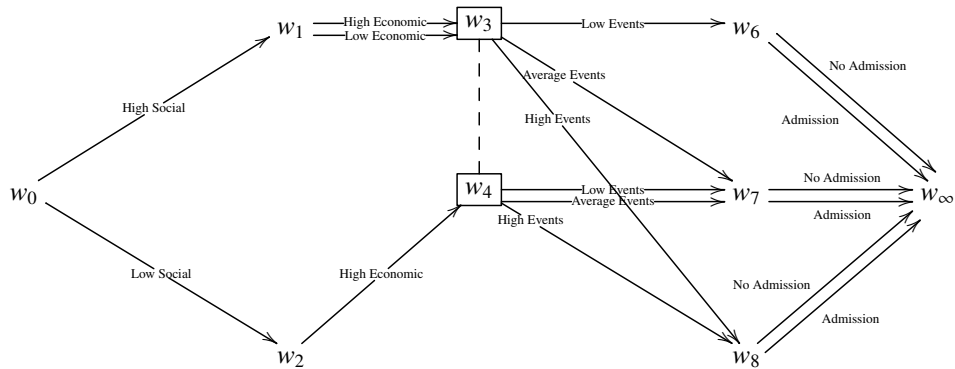


Figure 6: Manipulated CEG

The predicted probability of hospital admissions for families from a low social background and an average number of life events given this intervention is improved to 0.206 as can be found in Table 3. Further we can read from the table that the intervention reduces the predicted probability of a high number of life events in socially disadvantaged families from 0.468 to 0.196, giving an improvement in the probability of admissions mediated through the life events. We further emphasize here the point made by Shafer (1996) that causal assumptions should be inferred from tree-like structures as these naturally respect a causal ordering and in this sense we can see directly from the graph of the CEG and the associated conditional probability vectors at which position we may want to intervene.

6. Discussion

In this paper we have demonstrated through a new analysis that, whilst a BN search can be very useful for finding good explanations and providing a graphical framework for feeding back the analysis, the CEG provides useful embellishments to an initial BN study. This is not only apparent in the significantly high Bayes Factor of the derived CEG and the MAP BN structure but also in its expressiveness for the client.

There are several features of interest which could not be addressed here. In particular, the dynamic structure of the data could be used to provide further insight. Current research is being carried out on the development of a dynamic version of the CEG. The flexibility of the CEG of allowing the nodes of the probability tree to be combined in an asymmetric way can be exploited even more in the dynamic setting as nodes from different times points can be merged. With respect to health studies, such as the Christchurch Health and Development Study, these models would be very useful for describing the development of changing health status of the children due to family factors occurring at different points in time.

A further point of interest to be addressed in the future is the development of improved model selection algorithms to enable more efficient selection procedures. We are currently developing a dynamic programming algorithm for CEGs, based on Silander and Myllymäki (2006), which

is able to search the CEG space for problems with up to ten variables. Parallel to this the use of CEGs for more complex problems is being considered. The development of an ordinal CEG allows for an improved graphical representation by listing the positions in descending order according to a variable of interest. In addition to this it is often possible to reduce the original CEG structure by defining new variables resulting from the dependence structure of the variables depicted in the topology of the CEG. These aspects are discussed in detail in a subsequent paper (Barclay et al., 2012).

However, even when the analyses are further refined in the ways described above, the point illustrated in this paper still appears to hold: The embellishment of a BN into a CEG exhibits significant advantages in the application of graphical models to population studies like those of our example. We would therefore encourage BN modellers to explore the potential of enhancing a BN search with a subsequent CEG analysis on the full model or simply on a subset of the variables of the BN.

Acknowledgement

The authors would like to thank John Horwood and the CHDS research group for providing the data set.

References

- R. Cowell, A. Dawid, S. Lauritzen, D. Spiegelhalter, Probabilistic Networks and Expert Systems, Springer Verlag, 2007.
- K. Korb, A. Nicholson, Bayesian Artificial Intelligence, volume 1, cRc Press, 2004.
- D. Poole, N. Zhang, Exploiting contextual independence in probabilistic inference, *Journal of Artificial Intelligence Research* 18 (2003) 263–313.
- C. Bouilrier, N. Friedman, M. Goldszmidt, D. Koller, Context-specific independence in bayesian networks, in: *Proceedings of the Twelfth Conference on Uncertainty in Artificial Intelligence*, Portland, Oregon, pp. 115–123.
- N. Friedman, M. Goldszmidt, Learning bayesian networks with local structure, *Nata Asi Series D Behavioural and Social Sciences* 89 (1998) 421–460.
- J. Oliver, *Decision graphs-an extension of decision trees* (1993).
- M. Jaeger, Probabilistic decision graphs - combining verification and ai techniques for probabilistic inference, *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 12 (2004) 2004.
- J. Smith, P. Anderson, Conditional independence and chain event graphs, *Artificial Intelligence* 172 (2008) 42–68.
- G. Freeman, J. Smith, Bayesian map model selection of chain event graphs, *Journal of Multivariate Analysis* 102 (2011) 1152–1165.
- P. Thwaites, J. Smith, R. Cowell, Propagation using chain event graphs, in: *Proceedings of the Twenty-Fourth Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-08)*, AUAI Press, Corvallis, Oregon, 2008, pp. 546–553.
- P. Thwaites, J. Smith, E. Riccomagno, Causal analysis with chain event graphs, *Artificial Intelligence* 174 (2010) 889–909.
- P. Thwaites, Causal identifiability via chain event graphs, *Artificial Intelligence* (2012). To appear.
- Christchurch health and development study, <http://www.otago.ac.nz/christchurch/research/healthdevelopment/>, 2012. Accessed: Feb 2012.
- D. Fergusson, L. Horwood, F. Shannon, Social and family factors in childhood hospital admission., *Journal of Epidemiology and Community Health* 40 (1986) 50.
- D. Fergusson, M. Dimond, L. Horwood, F. Shannon, The utilisation of preschool health and education services, *Social Science & Medicine* 19 (1984) 1173–1180.
- S. Weich, J. Patterson, R. Shaw, S. Stewart-Brown, Family relationships in childhood and common psychiatric disorders in later life: systematic review of prospective studies, *The British Journal of Psychiatry* 194 (2009) 392–398.
- A. Beautrais, D. Fergusson, F. Shannon, Life events and childhood morbidity: A prospective study, *Pediatrics* 70 (1982) 935–940.

- D. Heckerman, D. Geiger, D. Chickering, Learning bayesian networks: The combination of knowledge and statistical data, *Machine learning* 20 (1995) 197–243.
- R. Neapolitan, *Learning Bayesian Networks*, Pearson Prentice Hall Upper Saddle River, NJ, 2004.
- S. Böttcher, C. Dethlefsen, deal: A package for learning bayesian networks, *Journal of Statistical Software* 8 (2003) 1–40.
- J. Pearl, *Causality: Models, Reasoning, and Inference*, 2nd edition, 2009.
- G. Shafer, *The Art of Causal Conjecture*, The MIT Press, 1996.
- T. Silander, P. Myllymäki, A simple approach for finding the globally optimal bayesian network structure, in: *Proceedings of the 22nd Conference on Artificial intelligence (UAI 2006)*, AUAI Press, 2006, pp. 445–452.
- L. Barclay, J. Hutton, J. Smith, Chain event graphs for informed missingness, *CRiSM Research Report 12-17* (2012).