Causality Enabled Compositional Modelling of Bayesian Networks

Jeroen Keppens and Qiang Shen
Department of Computer Science
University of Wales, Aberystwyth
{jrk,qqs}@aber.ac.uk

Abstract

Probabilistic abduction extends conventional symbolic abductive reasoning with Bayesian inference methods. This allows for the uncertainty underlying implications to be expressed with probabilities as well as assumptions, thus complementing the symbolic approach in situations where the use of a complete list of assumptions underlying inferences is not practical. However, probabilistic abduction has been of little use in first principle-based applications, such as abductive diagnosis, largely because no methods are available to automate the construction of probabilistic models, such as Bayesian networks (BNs). This paper addresses this issue by proposing a compositional modelling method for BNs.

Introduction

In many applications of model based diagnosis methods, it is generally assumed that the behaviour of a given system is deterministic and that there is only uncertainty with regards to which components behave abnormally and what the anomalous behaviour is. This may be a sensible assumption in the diagnosis of physical systems because there typically exists a good understanding of the way these systems work.

Recent research, presented in (Keppens & Zeleznikow 2003), has demonstrated that crime investigation can also be considered as a type of abductive diagnosis problem. However, the obvious distinction from the diagnosis of physical systems is that the behaviour of humans in a crime scenario is difficult to describe deterministically. There are too many assumptions that explain minor deviations between scenarios. An effective approach to overcoming this challenge is to model the behaviour in crime scenarios non-deterministically by integrating symbolic logic-based and probabilistic Bayesian methods (Pearl 1988; Poole 1993). On the one hand, the symbolic methods provide techniques for representing and reasoning with conditional independence relations, improving the efficiency of Bayesian inference. On the other hand, probabilistic methods can substitute the more obscure causes of uncertainty by probability distributions without having to explicitly recognise all of them.

In order for abductive diagnosis techniques to be effective, it is essential that the unknown scenarios and their behaviour can be composed from reusable parts (Coghill & Shen 2001; Struss & Price 2003). However, existing work on Bayesian Networks (BNs) assumes that the BNs are fully pre-specified, which may not be valid in solving may real-world problems. This paper introduces an approach to compositional modelling of BNs. This approach does not only determine the structure of BNs, but also the conditional probability distributions by composing the causal relations and probabilities attached to them. The method is illustrated with a simple abductive reasoning problem.

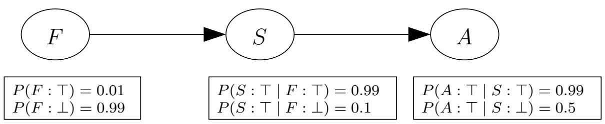

Figure 1: Sample Bayesian network

Towards Probabilistic Abductive Diagnosis

Inferring the plausible causes or explanations for known observations or symptoms by means of fixed domain knowledge is called abduction. For instance, if it is known that "fire" causes "smoke" and that "smoke" causes "alarm", then the observation of the alarm can be explained by fire. Developing an abductive reasoning system is difficult because formal domain knowledge is almost never entirely sound and complete. In the example, the rule smoke $\rightarrow$ alarm may be deemed unsound as it ignores various situations where smoke does not cause alarm, such as lack of a power supply and a malfunctioning sensor. Also, the two rules are obviously incomplete since there are many causes other than smoke, such as vapour and dirt, that are potential causes of alarm.

The lack of soundness and completeness can be represented by means of probabilities. Figure 1 represents a BN containing the domain knowledge of the smoke alarm example, extended with conditional probabilities. In this figure, the variables F, S and A respectively correspond to fire, smoke and alarm. The (potentially unsound) rule smoke $\rightarrow$ alarm is replaced by the conditional probability that alarm is true given that smoke is true $P(A:\top \mid S:\top) = 0.99$. As such, it allows for the 0.01 chance that smoke does not cause an alarm. Furthermore, the BN allows for the possibility that an alarm is triggered by causes other than smoke, as the conditional probability that alarm is true given that smoke is false $P(A:\top \mid S:\bot) = 0.5$.

Extending the abduction of the possible causes of symptoms or observations by calculating the likelihood of these causes is called probabilistic abduction (Poole 1993). In the ongoing example, for instance, the likelihood of fire as a cause of alarm can be computed as follows:

$$P(F:\top \mid A:\top) = \frac{\sum\limits_{i\in\{\top,\bot\}} P(A:\top \mid S:i)\times P(S:i \mid F:\top)\times P(F:\top)}{\sum\limits_{i\in\{\top,\bot\}} P(A:\top \mid S:i)\times \sum\limits_{j\in\{\top,\bot\}} P(S:i \mid F:j)\times P(F:j)}$$

One of the most important applications of abductive reasoning is abductive diagnosis (Console & Torasso 1990). An abductive diagnoser employs a knowledge base that contains combinations of causes that generate certain observation. Such knowledge is used to abductively infer plausible scenarios that may explain the known symptoms, and to determine possible source of evidence that can help reduce the set of plausible scenarios.

The knowledge base of an abductive diagnoser normally follows the compositional modelling paradigm (Falkenhainer & Forbus 1991; Keppens & Shen 2001): Its knowledge is stored in the form of reusable parts of scenarios that can be composed to meet the specific circumstances of a given problem. In the conventional application domain of diagnosis of physical systems, these parts are model fragments containing differential equations describing the behaviour of components in a certain configuration. Most compositional modellers can combine such model fragments to form models whose simulated behaviour matches the observations (Heller & Struss 1998; 2001; Levy, Iwasaki, & Fikes 1997; Nayak & Joskowicz 1996; Rickel & Porter 1997). In more recent work (Keppens & Zeleznikow 2003), the same approach has also been applied to descriptions in predicate logic of the events leading up to crime evidence.

There are obviously substantial benefits to integrating probabilistic abduction into abductive diagnosis. Probabilities can express uncertainty in the causal relations and help enable the use of maximal expected entropy reduction techniques (de Kleer & Williams 1987) to suggest evidence collection strategies. As has been argued in the literature (Coghill & Shen 2001; Struss & Price 2003), this is best achieved in conjunction with some means to compose a space of plausible models, represented as a BN in this case, from observed symptoms. The next two sections devise such a compositional modeller.

Knowledge Representation

Before explaining the proposed approach to automated composition of BNs from knowledge, the knowledge representation formalism adopted and its underlying assumptions are discussed.

Intuitions

The purpose of having a knowledge representation formalism is to describe scenarios and their relations to evidence or symptoms in a concise manner. Scenarios can be seen to consist of states and events. A *state* is a particular status of the universe of discourse, often specific to a particular instance in time and space. Typical examples of a state are "there is smoke in the room", "5 red cotton fibres have been retrieved from John's black woollen jumper", and "a small amount of skin tissue is present under Jane's fingernails". An *event* is a phenomenon that creates a new state or causes a change to the state of a part of the universe of discourse. Typical examples of events are "a fire generates smoke", "5 fibres are transferred from Jane's red cotton t-shirt to John's black woollen jumper", and "Jane scratches John with her fingernails in self-defence".

In order to use such scenarios for diagnosis, the events and states must also be related to one another. Although there are many ways to link the constituents of a scenario, for simplicity this work will only use one type of representation to describe the valid forms of inference: *causal relations*. As to be shown in the next section, this representation scheme helps automated composition of BNs considerably. In particular, the use of causal relations gives rise to the underlying presumption of the present work in that each scenario constituent is either *assumed* to be true, or *triggered* as a consequence of a combination of other states and events. This is reasonable in many application domains, including the descriptions of crime scenarios.

Formal Notation

Within graphical models that describe probabilistic dependencies in general and BNs in particular, the information of interest, here the states and events of plausible scenarios, is represented by means of variables with discrete domains. Although it may be set by default that these will be truth/falsehood assignments to predicates, this is often insufficient. In many diagnostic applications several types of event may each contribute to or detract from a particular state. In the ongoing smoke alarm example, for instance, it is the combined volume of smoke and vapour particles that determines the likelihood of the activation of the smoke alarm, rather than the constituent causes. In such cases, it is necessary to represent some states and events quantitatively or semi-quantitatively.

Having taken notice of this, states and events are herein represented as *variable assignments*. Variables are defined as tuples $\langle p, D_p\rangle$ where p is a predicate identifying the variable, and D_p is a domain of values. Unless otherwise specified, variables are assumed to have the boolean domain $\{\top, bot\}$, although other domains are allowed. Typical examples of variable definitions are:

$$\langle \text{smoke}, \{\top,\bot\}\rangle$$
$$\langle \text{transfer(fibres,john,jane)}, \mathbb{N}\rangle$$
$$\langle \text{transfer(skin-tissue,john,jane)}, \{\text{some}, \text{infinitesimal}, \text{some}\}\rangle$$

A *variable assignment* is any constraint on the domain of a variable. Generally speaking, the assignment of a set of values D' to a variable identified by predicate p, where $D' \subseteq D_p$, is denoted by $p : D'$. In practice, however, more useful expressions for defining variable assignments may be employed, such as the assignment of an individual value:

$$p : v \quad \text{where } v \in D_p$$

and, if the variable has an ordered domain, the assignment of a lower and/or an upper boundary:

$$v_l < p < v_h \quad \text{where } v_l, v_h \in D_p$$

Causal relations between variable assignments are expressions of the form:

$$p_1 : D'_{p_1} \wedge \ldots \wedge p_m : D'_{p_m}$$
$$\xrightarrow[\oplus]{C} p_n : \{f(v_1,\ldots,v_m) \mid v_1 \in D'_{p_1}, \ldots, v_m \in D'_{p_m}\}$$

where

- f is a function $D_{p_1} \times \ldots \times D_{p_m} \mapsto D_{p_n}$. This function computes the assignment of the consequent variable p_n on the basis of the assignments of the antecedent variables $p_1, \ldots, p_m$.
- $\oplus$ is a commutative and associative function $D_{p_n} \times D_{p_n} \mapsto D_{p_n}$. It is the *composition operator*, describing how the effects of different causes must be combined. By default, this operator is taken to be the logical disjunction operator $\vee$, implying that a single possible cause is sufficient for the consequent state or event to be true.

Such causal relations are considered as q*influences*, whereby a combination of states and events described by the assignment of antecedent variables results in a state or event denoted by the assignment of the consequent variable. In other words, they are similar to the influences in QPT (Forbus 1984), but adapted to denote non-physical processes. As in QPT, multiple influences may independently affect the consequent state/event. The composition operator defines how the outcomes of different influences are combined. The meaning and usage of causal relations is illustrated by means of the following examples: The causal relation

$$\text{smoke} : \top \xrightarrow[\vee]{C} \text{alarm} : \top \quad (1)$$

states that smoke triggers the (smoke) alarm. The composition operator $\vee$ is employed because smoke is in itself a sufficient condition to trigger the alarm. The causal relation

$$\text{strangles}(P1, P2) : \top \xrightarrow[+]{C} \text{transfer}(\text{fibres}, P1, P2) > 10$$

states that if person $P1$ strangles person $P2$ at least 10 fibres are transferred from (the clothes of) $P1$ to (the clothes of) $P2$ as a consequence of that event. The combination operator $+$ (algebraic sum) is employed because other events may cause additional transfers of fibres, adding to the total number of transferred fibres.

Knowledge Base

The knowledge base employed herein consists of prior probability distributions for the variables describing *assumed* states and events, and so-called *probabilistic causal relations* (PCRs) that explain the values of the variables corresponding to *triggered* states and events. PCRs are causal relations that may have any one of a set of possible outcomes as its consequence. This is clearly distinct from conventional symbolic causal relations, which have a single predetermined consequence. Which of the multiple possible outcomes is the actual consequence of a given PCR depends on a probability distribution.

Formally, a PCR is a tuple $\langle P_a, D_a, p_c, \oplus, f\rangle$, where:

- $P_a = \{p_1, \ldots, p_m\}$ is a set of participants, called the *antecedent participants*;
- $D_a \subset D_{p_1} \times \ldots \times D_{p_m}$ is a subset of the Cartesian product of the domains of the antecedent participant;
- p_c is a participant, called the *consequent participant*;
- $\oplus$ is the composition operator over the causal relations; and
- f is a function $D_a \times D_{p_c} \mapsto [0, 1]$ that specifies the probability of obtaining an assignment $p_c : v_c$, with $v_c \in D_{p_c}$, as a consequence of this PCR, given an antecedent p_c : $v_1, \ldots, p_m : v_m$, with $(v_1, \ldots, v_m) \in D_a$. The latter probability of the actual outcome of a PCR is denoted by

$$P(p_1 : v_1 \wedge \ldots \wedge p_m : v_m \xrightarrow[\oplus]{C} p_c : v_c) = f(v_1, \ldots, v_m, v_c) \quad (2)$$

For example, the PCR

$$\langle\{\text{smoke}\}, \{(\top)\}, \text{alarm}, \vee, f\rangle$$
$$\text{with } f(\top, \top) = 0.99 \text{ and } f(\top, \bot) = 0.01$$

is a probabilistic version of the causal relation described in (1). It states that there is a 99% chance that smoke triggers the alarm and a 1% probability that it does not.

To enable their use in compositional modelling of BNs, it is presumed that the PCRs in a given knowledge base possess the following properties:

1. *All PCRs with the same consequent participant employ the same composition operator* $\oplus$. For notational convenience, a single composition operator is assumed to describe how the outcomes of different causal relations are combined to represent their aggregate effect on a single consequent participant. However, when necessary, this assumption can be relaxed to allow for multiple composition operators to be used, as long as they can be composed (Bobrow *et al.* 1996; Keppens & Shen 2004).
2. *The PCRs in a knowledge base do not contain functions that map the same combination of assignments of the antecedent participants to the same assignment of the consequent participant.* In other words, for any given set of (antecedent) variable assignments $A = \{p_1 : v_1, \ldots, p_m : v_m\}$ and any (consequent) participant assignment $p_c : v_c$, there is at most one PCR $\langle P_a, D_a, p_c, \oplus, f\rangle$ such that
$$P_a = \{p_1, \ldots, p_m\},$$
$$(v_1, \ldots, v_m) \in D_a, \text{ and}$$
$$f(v_1, \ldots, v_m, v_c) > 0$$
This assumption facilitates the specification of probability distributions over the multiple possible outcomes of the PCR, as described in the next assumption.
3. *All PCRs with the same antecedent and consequent participants define a probability distribution.* Let P_a be the set of participants $\{p_1, \ldots, p_m\}$ and $A = \{p_1 : v_1, \ldots, p_m : v_m\}$ be a set of variable assignments to the participants of P_a such that one or more PCRs of the form $\langle P_a, D_a, p_c, \oplus, f\rangle$ exist in a knowledge base, with $(v_1, \ldots, v_m) \in D_a$. Then, the function $\gamma_{A,p_c} : D_{p_c} \mapsto [0, 1]$ with
$$\gamma_{A,p_c}(v_c) = \begin{cases} f(v_1, \ldots, v_m, v_c) & \text{if a PCR } \langle P_a, D_a, p_c, \oplus, f\rangle \\ & \text{exists with } (v_1, \ldots, v_m) \in D_a \\ & \text{and } f(v_1, \ldots, v_m, v_c) > 0; \\ 0 & \text{otherwise.} \end{cases}$$
is assumed to be a probability distribution. γ_{A,p_c} is called the probability distribution defined by the knowledge base with respect to the set of assignments A and consequent participant p_c. This assumption formalises the concept that a PCR is a causal relation with a non-deterministic outcome governed by a probability distribution.

4. *Every pair of probability distributions defined by the PCRs with respect to the same consequent participant are independent.* Intuitively, this assumption implies that the outcome of an influence underlying one PCR affecting the value of a variable p_c is not affected by that of another. This assumption facilitates the computation of a combined effect of distinct causal relations. Let there be two influences affecting the value of p_c, one whose outcome is dependent on the values of the variables in P_1 and one dependent on those in P_2, and let A_1 and A_2 be sets of assignments to the variables in P_1 and P_2 respectively. Then, the probability that the first influence yields $p_c : v_1$ and that the second yields $p_c : v_2$ can be computed by multiplying $\gamma_{A_1,p_c}(v_1)$ and $\gamma_{A_2,p_c}(v_2)$, since it follows from the independence assumption that:

$$P\left(\left[\left(\wedge_{a_{1i}\in A_1} a_{1i}\right) \overset{C}{\underset{\oplus}{\rightarrow}} p_c : v_1\right] \wedge \left[\left(\wedge_{a_{2i}\in A_2} a_{2i}\right) \overset{C}{\underset{\oplus}{\rightarrow}} p_c : v_2\right]\right)$$
$$= P\left[\left(\wedge_{a_{1i}\in A_1} a_{1i}\right) \overset{C}{\underset{\oplus}{\rightarrow}} p_c : v_1\right] \times P\left[\left(\wedge_{a_{2i}\in A_2} a_{2i}\right) \overset{C}{\underset{\oplus}{\rightarrow}} p_c : v_2\right]$$
$$= \gamma_{A_1,p_c}(v_1) \times \gamma_{A_2,p_c}(v_2)$$

5. *There are no cycles in the knowledge base*. This means that there is no subset of PCRs in the knowledge base which is of the form:

$$\begin{gathered}\langle\{\ldots,p_1,\ldots\}, D_{a1}, p_2, \oplus_1, f_1\rangle\\ \langle\{\ldots,p_2,\ldots\}, D_{a2}, p_3, \oplus_2, f_2\rangle\\ \vdots\\ \langle\{\ldots,p_n,\ldots\}, D_{an}, p_1, \oplus_n, f_n\rangle\end{gathered}$$

This assumption is required because BNs can not represent such information as they are inherently acyclic (Pearl 1988).

Note that assumptions 1 and 2 above are introduced merely for notational convenience. Assumption 3 is inherent to the notion of PCRs as it imposes the requirements of probability theory upon the measurements of likelihood of the potential outcomes of the PCRs. At first sight, assumption 4 may seem controversial. However, by returning to the original intuition that causal relations represent underlying influences, it becomes clear how a knowledge base can be engineered to satisfy this assumption. Two probability distributions associated with PCRs that share the same consequent participant may be dependent for the following reasons:

- *The underlying influences share a sub-influence*. This situation can be resolved by redefining one of the PCRs such that the shared influence is removed from the influence underlying it. Consider, for example, one PCR describing that fire causes smoke and another describing that cigarettes cause smoke. Clearly, the probability distribution of the amount of smoke caused by cigarettes is not independent from the probability distribution of the amount of smoke caused by fire, because cigarettes may cause fire. Yet, if the PCR containing the causal link between cigarettes and smoke counts only the smoke produced directly by the cigarettes whilst excluding the smoke generated by a possible fire that may have been stared by the same cigarette, then it becomes feasible to define independent probability distributions.
- *There are latent variables affecting both influences*. In the previous example, a small dependency may still appear to remain between the redefined PCRs because the amount of smoke generated by cigarettes of that by fire varies with respect to the composition of the air. By including the influences of the latent variables affecting both processes, e.g. the proportion of oxygen in the air, remaining dependencies can then be explained away.

Finally, assumption 5 prohibits the use of feedback loops in the knowledge base to allow the construction of BNs. In the chosen application domain of crime scenarios, this is a reasonable assumption because evidence is normally the result of a finite sequence of events. Existing work on the development of models of the causes of evidence, such as (Aitken, Taroni, & Garbolino 2003; Evett *et al.* 2000; Mortera, Dawid, & Lauritzen 2003), supports this hypothesis as it makes extensive use of (manually constructed) BNs.

Note that although PCRs describe influences between variables, they differ from the influences in QPT in three aspects. Firstly, the outcomes of PCRs are probabilistic as opposed to the deterministic outcomes of QPT influences. Secondly, the representation scheme of PCRs is more extensive to allow for a wider variety of variable domains and composition operators. Thirdly, a knowledge base of PCRs must not contain cycles of influences because BNs will be employed to infer information from them.

Bayesian Networks Composition

As BNs consist of two distinct features, a directed acyclic graph (DAG) and a set of conditional probability tables, this section is divided into two subsections describing how both aspects can be composed automatically from a given knowledge base.

Structural Composition

Automated composition of the structure of a BN from partial dependency relations is not a trivial task because the structure must uphold the semantics of the BN paradigm. Suppose that G is a graph that is constructed by the following procedure:

- *Initialisation*: The original part of G is first created with one node for each piece of evidence. Each of these nodes contains a variable p describing the type of evidence it corresponds to.
- *Backward chaining phase*: For each PCR $\langle P_a, D_a, p_c, \oplus, f\rangle$, such that p_c corresponds to a node already in G, a node is added for each participant in P_a for which no node already exists, and an arc is added from p_a to p_c.
- *Forward chaining phase*: For each PCR $\langle P_a, D_a, p_c, \oplus, f\rangle$, such that the variables in P_a correspond to nodes already in G, a node corresponding to p_c is added to G if one does not already exist and arcs from each node corresponding to a variable in P_a to p_c are added.

It can then be shown that G has two important properties as follows. First, G is a directed acyclic graph (DAG). This property follows directly from assumption 5 of the knowledge base. Second, G is a minimal independence map of the conditional independence model, a formal requirement of any DAG being a BN (Pearl 1988).

Here, a conditional independence model describes for each pair of variables p_x and p_y and for each set of variables P, excluding p_x and p_y, whether information on the value of p_y affects the probability of values of p_x if the values of the variables in P are already known. Two variables p_x and p_y are said to be conditionally independent given a set of variables $P = \{p_i, \ldots, p_j\}$ if

$$P(p_x : v_x \mid p_i : v_i \wedge \ldots \wedge p_j : v_j \wedge p_y : v_y) \\ = P(p_x : v_x \mid p_i : v_i \wedge \ldots \wedge p_j : v_j)$$

A formal proof that the DAG G does satisfy this property is beyond the scope of this paper. However, it is worth noting that it can be shown that the requirements for two variables to be conditionally independent given a set of variables P, with respect to a knowledge base of PCRs, are equivalent to the requirements for those two variables to be d-separated by P in the DAG G. Following (Pearl 1988), this implies that G is a minimal independence map of the conditional independence model represented by the PCRs.

Parameter Composition

A BN also requires a complete specification of the conditional probability tables to be of any practical use. Let m be the number of states of each node in the BN and q be the number of parents of each non-root node. Then, a total of $m^q \times (m-1)$ probabilities must be assigned to each non-root node. In an abductive diagnosis application, q may become large, thus inhibiting the manual specification of the conditional probabilities.

Using the method to derive the structure of a BN from knowledge, which is described in the previous subsection, it is possible to determine the set of (immediate) parent variables $P = \{p_1, \ldots, p_m\}$ of a variable p_c. The following explains how the adopted knowledge representation scheme can be employed to compute conditional probabilities $P(p_c : v_c \mid A)$, where A is a set of assignments $\{p_1 : v_{p_1}, \ldots, p_m : d_{v_m}\}$.

As explained earlier, the value of p_c is determined by the values of the variables in P and the effect of the PCRs from variables in P upon p_c. Any combination of causal relations, whose antecedent is entailed by the set of assignment A and whose composition of consequences (as computed with the composition operator) equals the value v_c, is sufficient to explain $p_c : v_c$. Thus,

$$P(p_c : v_c \mid A) = \\ P\Bigg(\bigvee_{\substack{[v_1 \oplus \ldots \oplus v_n = v_c], \\ \{A_1, \ldots, A_n\} = \pi(A)}} \bigg(\bigwedge_{i=1,\ldots,n} (A_i \overset{C}{\underset{\oplus}{\rightarrow}} p_c : v_i) \bigg) \Bigg)$$

where

- $\pi(A)$ is the (complete) set $\{A_1, \ldots, A_n\}$ of all subsets of A that correspond to the assignments in the antecedent of a PCR that affects the value of p_c, and
- $\{v_1, \ldots, v_n\}$ is a set of plausible outcomes of the n relevant PCRs, under the n corresponding assignments $A_1, \ldots, A_n$, such that the combined effect of all PCRs $v_1 \oplus \ldots \oplus v_n$ equals v_c.

The above equation states that $P(p_c : v_c \mid A)$ equals the probability of a disjunctive normal form (DNF). The conjunctions within this DNF describe situations where n PCRs yield outcomes $p_c : v_1, \ldots, p_c : v_n$, with the combination of these outcomes $v_1 \oplus \ldots \oplus v_n$ equalling v_c. The DNF considers the disjunction of all possible conjunctions of outcomes $p_c : v_1, \ldots, p_c : v_n$ with $v_1 \oplus \ldots \oplus v_n = v_c$.

Because a single PCR can not yield two outcomes simultaneously under any given assignments of the antecedent participants, the different disjuncts in the DNF are mutually exclusive. In other words, the probability of two disjuncts occurring at the same time is 0. Therefore, $P(p_c : v_c \mid A)$ equals the sum of the probabilities of all the disjuncts:

$$P(p_c : v_c \mid A) = \sum_{\substack{[v_1 \oplus \ldots \oplus v_n = v_c], \\ \{A_1, \ldots, A_n\} = \pi(A)}} P\Bigg(\bigwedge_{i=1,\ldots,n} (A_i \overset{C}{\underset{\oplus}{\rightarrow}} p_c : v_i) \Bigg)$$

According to assumption 4 made previously regarding the knowledge base, the probability distributions defined by the PCRs with respect to the same consequent participant are independent. Therefore, the probability of a conjunction of the outcomes of certain PCRs with the same consequent participant is given by the product of the probabilities of those outcomes:

$$P(p_c : v_c \mid A) = \sum_{\substack{[v_1 \oplus \ldots \oplus v_n = v_c], \\ \{A_1, \ldots, A_n\} = \pi(A)}} \prod_{i=1,\ldots,n} P(A_i \overset{C}{\underset{\oplus}{\rightarrow}} p_c : v_i)$$

The above equation can also be expressed in terms of the probability distributions γ_{A_i, p_c}. Let Γ_{A,p_c} be the set of all probability distributions defined by the PCRs with respect to a subset of the variable assignments in A and the consequent participant p_c. Then,

$$\Gamma_{A,p_c} = \{\gamma_{A_i,p_c} \mid A_i \subseteq A\}$$

If $\Gamma_{A,p_c} = \{\gamma_{A_1,p_c}, \ldots, \gamma_{A_n,p_c}\}$, where n is the number of probability distributions defined by the PCRs with respect to a subset of A and consequent participant p_c, then

$$P(p_c : d_c \mid A) = \sum_{\substack{d_1 \oplus \ldots \oplus d_n = d_c \\ n = |\Gamma_{A,p_c}|}} \Bigg(\prod_{i=1,\ldots,n} \gamma_{A_i,p_c}(d_i) \Bigg) \quad (3)$$

In the next section, it will be shown how (3) can be applied to a given probabilistic abductive reasoning problem.

An Illustrative Example

To illustrate the ideas of this work, the following discusses a knowledge base for the aforementioned smoke alarm diagnosis problem and shows how this knowledge base can be used to construct a BN for probabilistic abductive diagnosis.

The knowledge base of model fragments is constructed by considering the influences that may affect the types of evidence of interest and/or by considering the types of scenario that the knowledge base may have to deal with and their consequences. In the smoke alarm example, situations that trigger the alarm are examined. Smoke alarms function by sensing certain types of particles in the air, and hence, they may be activated by smoke, dust and vapour particle[1]. Common events influencing the amount of smoke in a room are smoking cigarettes, fire and cooking. Cooking may also generate vapours. Below, a knowledge base that formalises this intuitive knowledge is given. In application domains where such knowledge is available in terms

[1]This example aims at demonstrating how a variety of situations can be considered without being overly complex. Although there are clearly other possible types of particle that may trigger a smoke alarm, they are not considered here to limit the size of this discussion.

Name	Description	Domain
alarm	smoke alarm sounds	B
cigarettes	cigarettes are smoked near alarm (without causing fire)	B
cooking	someone is cooking near alarm	B
dust	amount of dust near alarm	Q
fire	there is a fire near alarm	B
smoke	amount of smoke near alarm	Q
vapours	amount of vapours near alarm	Q

Table 1: Variables in the knowledge base

$P(\text{cigarettes} : \top) = 10.00\%,\ P(\text{cigarettes} : \bot) = 90.00\%$
$P(\text{cooking} : \top) = 10.00\%,\ P(\text{cooking} : \bot) = 90.00\%$
$P(\text{dust} : \text{none}) = 10.00\%,\ P(\text{dust} : \text{some}) = 80.00\%,$
$P(\text{dust} : \text{much}) = 10.00\%$
$P(\text{fire} : \top) = 1.00\%,\ P(\text{fire} : \bot) = 99.00\%$

Table 2: Prior probabilities

of expertise, the construction of the knowledge base may be expected to be carried out via a conventional knowledge acquisition process.

Table 1 lists the 7 variables that are considered in the knowledge base, as well as a description of their meaning and domain. Two domains are employed in this example:

$$B = \{\top, \bot\}$$
$$Q = \{\text{none}, \text{some}, \text{many}\}$$

B is the conventional Boolean domain, where $\top$ corresponds to true and $\bot$ to false. The logical disjunction operator $\vee$ is defined over this domain and is used as the composition operator for probabilistic causal relations with a Boolean consequent variable. Q is an ordered set containing qualitative expressions of quantities, with none $<$ some $<$ many. The $\max$ operator, which returns the highest value of a pair of values taken from Q, is used as a composition operator for probabilistic causal relations with a consequent variable that has Q as its domain.

B and Q are chosen here to demonstrate the use of different domains and composition operators, whilst keeping the presentation of this example manageable. Obviously, in a more extensive application, more complex domains could be employed. For instance, instead of Q, a finite partition of the domain of real numbers $\mathbb{R}$ may be used, expressing the number of particles of dust, smoke and vapour per m^3.

The knowledge base consists of a set of prior probabilities for variables whose values are not affected causally and a set of probabilistic causal relations (PCRs). The structure of the PCRs corresponds to the influences acquired from the experts. In the literature on forensic statistics, the probabilities are typically subjective ones provided by experts (Aitken, Taroni, & Garbolino 2003; Cook *et al.* 1999). The prior probabilities are listed in Table 2 and the content of the PCRs is described in Tables 3, 4 and 5. The PCRs in this example all have one antecedent variable and one consequent variable. In the three tables summarising them, the rows correspond to the assignment of the antecedent variable, the columns to the assignment of the consequent variables and the cells denote the probability of the causal relations. For example, the entry 29.00% in column 2, row 3 of Table 3 states that:

$$P(\text{cigarettes} : \top \xrightarrow[\max]{C} \text{smoke} : \text{none}) = 0.29$$

Antecedent	smoke		
	none	some	much
cigarettes:$\top$	29.00%	70.00%	1.00%
cigarettes:$\bot$	100.00%	0.00%	0.00%
fire:$\top$	1.00%	29.00%	70.00%
fire:$\bot$	100.00%	0.00%	0.00%
cooking:$\top$	80.00%	19.00%	1.00%
cooking:$\bot$	100.00%	0.00%	0.00%

Table 3: Causal probabilities of the form $P(\text{Antecedent} \xrightarrow[\max]{C} \text{smoke} : v_s)$

Antecedent	vapours		
	none	some	much
cooking:$\top$	10.00%	70.00%	20.00%
cooking:$\bot$	100.00%	0.00%	0.00%

Table 4: Causal probabilities of the form $P(\text{Antecedent} \xrightarrow[\max]{C} \text{vapours} : v_v)$

A BN can be derived from these prior probabilities and PCRs. In particular, the structure is determined by the form of the PCRs. For each of the 7 variables in these relations, a node is created in the BN. And for each PCR based on a causal relation of the form $p_i : v_i \xrightarrow[\oplus]{C} p_j : v_j$, where p_i and p_j are predicates describing variables, v_i and v_j are values in the domains of p_i and p_j respectively, and $\oplus$ is either $\vee$ or $\max$, an arc is drawn from the node that corresponds to p_i to the one associated with p_j. The structure of the resulting BN is shown in Figure 2.

The conditional probability tables are computed from the probabilities of the causal relations. For instance, the PCRs "cigarettes", "fire" and "cooking" to "smoke" are combined as follows:

$$P(\text{smoke} : v_s \mid \text{cigarettes} : v_{ci}, \text{fire} : v_f, \text{cooking} : v_{co}) = \sum_{\max(v_{s,ci}, v_{s,f}, v_{s,co}) = v_s} \begin{pmatrix} P(\text{cigarettes} : v_{ci} \xrightarrow[\max]{C} \text{smoke} : v_{s,ci}) \times \\ P(\text{fire} : v_f \xrightarrow[\max]{C} \text{smoke} : v_{s,f}) \times \\ P(\text{cooking} : v_{co} \xrightarrow[\max]{C} \text{smoke} : v_{s,co}) \end{pmatrix} \quad (4)$$

Antecedent	alarm	
	$\top$	$\bot$
dust:none	0.00%	100.00%
dust:some	1.00%	99.00%
dust:much	3.00%	97.00%
smoke:none	0.00%	100.00%
smoke:some	20.00%	80.00%
smoke:much	99.00%	1.00%
vapours:none	0.00%	100.00%
vapours:some	1.00%	99.00%
vapours:much	5.00%	95.00%

Table 5: Causal probabilities of the form $P(\text{Antecedent} \xrightarrow[\vee]{C} \text{alarm} : v_a)$

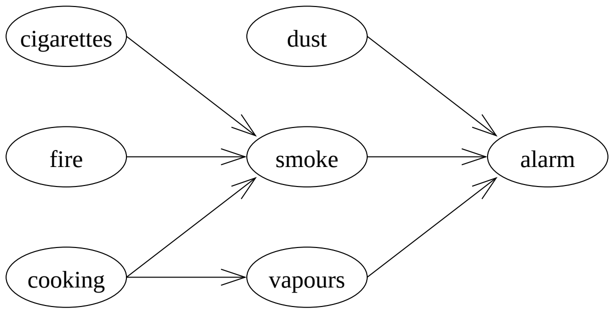

Figure 2: Bayesian Network Structure

cigarettes	fire	cooking	smoke		
			none	some	much
⊤	⊤	⊤	0.23%	29.17%	70.60%
		⊥	0.29%	29.41%	70.30%
	⊥	⊤	23.20%	74.81%	1.99%
		⊥	29.00%	70.00%	1.00%
⊥	⊤	⊤	0.80%	28.90%	70.30%
		⊥	1.00%	29.00%	70.00%
	⊥	⊤	80.00%	19.00%	1.00%
		⊥	100.0%	0.00%	0.00%

Table 6: Conditional probabilities of the form $P(\text{smoke} : v_s \mid \text{cigarettes} : v_{ci}, \text{fire} : v_f, \text{cooking} : v_{co})$

To calculate the conditional probability of smoke : v_s given cigarettes : v_{ci}, fire : v_f, and cooking : v_{co}, all combinations of causal relations (cigarettes : $v_{ci} \overset{C}{\underset{\max}{\rightarrow}}$ smoke : $v_{s,ci}$), (fire : $v_f \overset{C}{\underset{\max}{\rightarrow}}$ smoke : $v_{s,f}$), and (cooking : $v_{co} \overset{C}{\underset{\max}{\rightarrow}}$ smoke : $v_{s,co}$) are considered, such that the maximum of $v_{s,ci}$, $v_{s,f}$, and $v_{s,co}$ equals v_s. It computes the probability of the conjunction of causal relations in each combination by taking their product and the probability of the disjunction of all combinations by summing the products. The following two examples demonstrate the computation of the conditional probabilities of smoke : none and smoke : some given cigarettes : ⊤, fire : ⊤ and cooking : ⊤:

$$\begin{aligned}
&P(\text{smoke : none} \mid \text{cigarettes} : \top, \text{fire} : \top, \text{cooking} : \top)\\
=&P(\text{cigarettes} : \top \overset{C}{\underset{\max}{\rightarrow}} \text{smoke : none})\\
&\times P(\text{fire} : \top \overset{C}{\underset{\max}{\rightarrow}} \text{smoke : none})\\
&\times P(\text{cooking} : \top \overset{C}{\underset{\max}{\rightarrow}} \text{smoke : none})\\
=&0.29 \times 0.01 \times 0.80 = 0.00232 \text{ (or } 0.23\%)
\end{aligned}$$

$$\begin{aligned}
&P(\text{smoke : some} \mid \text{cigarettes} : \top, \text{fire} : \top, \text{cooking} : \top)\\
=&(0.29 \times 0.01 \times 0.19) + (0.29 \times 0.29 \times 0.80) +\\
&(0.29 \times 0.29 \times 0.19) + (0.70 \times 0.01 \times 0.80) +\\
&(0.70 \times 0.01 \times 0.19) + (0.70 \times 0.29 \times 0.80) +\\
&(0.70 \times 0.29 \times 0.19) = 0.29171 \text{ (or } 29.17\%)
\end{aligned}$$

Tables 6, 7 and 8 show the conditional probabilities computed from the probabilistic causal relations in Tables 3, 4 and 5 respectively. The DAG in Figure 2 and these resultant conditional probabilities completely specify a BN.

cooking	vapours		
	none	some	much
⊤	10.00%	70.00%	20.00%
⊥	100.00%	0.00%	0.00%

Table 7: Conditional probabilities of the form $P(\text{vapours} : v_v \mid \text{cooking} : v_{co})$

dust	smoke	vapours	alarm	
			⊤	⊥
none	none	none	0.00%	100.00%
		some	1.00%	99.00%
		much	5.00%	95.00%
	some	none	20.00%	80.00%
		some	20.80%	79.20%
		much	24.00%	76.00%
	much	none	99.00%	1.00%
		some	99.01%	0.99%
		much	99.05%	0.95%
some	none	none	1.00%	99.00%
		some	1.99%	98.01%
		much	5.95%	94.05%
	some	none	20.80%	79.20%
		some	21.59%	78.41%
		much	24.76%	75.24%
	much	none	99.01%	0.99%
		some	99.02%	0.98%
		much	99.06%	0.94%
much	none	none	3.00%	97.00%
		some	3.97%	96.03%
		much	7.85%	92.15%
	some	none	22.40%	77.60%
		some	23.18%	76.82%
		much	26.28%	73.72%
	much	none	99.03%	0.97%
		some	99.04%	0.96%
		much	99.08%	0.92%

Table 8: Conditional probabilities of the form $P(\text{alarm} : v_a \mid \text{dust} : v_d, \text{smoke} : v_s, \text{vapours} : v_v)$

The composed BN can be employed for various diagnostic analyses. Consider, for example, the problem of determining whether or not the smoke alarm was triggered by a fire or some other cause. The conventional Bayesian approach for testing such a hypothesis would be to compute the likelihood ration using the information stored in the BN:

$$L(\text{alarm} \mid \text{fire}) = \frac{P(\text{alarm} \mid \text{fire})}{P(\text{alarm} \mid \neg\text{fire})} = \frac{0.7548}{0.0319} = 23.46$$

If the knowledge base, and the corresponding BN, would also contain other potential investigative actions and corresponding evidence, maximum expected entropy reduction techniques could be applied to devise efficient evidence collection strategies. However, this remains a piece of future research.

Conclusion and Future Work

In this paper a compositional modelling approach to construct Bayesian networks (BNs) from reusable parts, representing probabilistic causal relations, has been presented. By means of a knowledge base of causal relations, it composes DAGs that satisfy the semantics of BN structures. And by attaching probabilities to the causal relations, the approach also allows for the composition of conditional probability tables, thereby constructing fully specified BNs.

As such, this work has addressed an important limitation of BNs. Previously, BNs had to be completely specified (or learned from data) by the knowledge engineer (Pearl 1988). To test different hypotheses, either different BNs had to be specified completely, or a single more extensive BN had to be created (which increases the size of the conditional probability tables, and hence the computational complexity). Compositional modelling of BNs enables the flexible adaptation of probabilistic graphical models to different problems and situations.

The primary motivation of this work has been to integrate the approach into a probabilistic abductive diagnostic systems. The main application area of this work is currently set to automate the evaluation of forensic evidence. As explained in (Keppens & Zeleznikow 2003), this forms a diagnostic task that may be best resolved abductively. However, there are too many plausible assumptions underlying the availability or lack of forensic evidence to list and evaluate them explicitly. For this reason, forensic scientists have suggested the use of BNs to analyse the implications of forensic evidence and the expected benefit of evidence collection strategies (Evett *et al.* 2000). This work helps the development of a decision support system towards automating such analyses.

The next logical step in this work is the development of a knowledge base in a real-world domain to evaluate its scaleupability. The forensic statistics literature has devised a substantial number of BN instances for the analysis of forensic evidence, such as DNA mixtures (Mortera, Dawid, & Lauritzen 2003) and transferred blood (Aitken, Taroni, & Garbolino 2003). From such BNs, the constituent component parts can be extracted and formalised as probabilistic causal relations in a knowledge base upon which to perform the evaluation.

Another issue that is beyond the scope of this paper is the estimation of probabilities of both the priors and the given probabilistic causal relations. Some of these probabilities may be determined by experimentation whilst others can be purely subjective. Because it is often inappropriate to express subjective probabilities with precise numbers, semi-quantitative methods, such as linguistic probabilities (Halliwell & Shen 2002), can provide an effective means of representing them in domains such as forensic statistics. Therefore, integrating semi-quantitative representations of probability into the approach described herein remains as an important piece of further research.

Acknowledgements

This work has been partly supported by UK EPSRC grant GR/S63267 and much of it was carried out while the authors were with the School of Informatics, the University of Edinburgh. The authors are grateful to Colin Aitken, Joe Halliwell and Burkhard Schafer for their contribution, whilst taking full responsibility for the views expressed in this paper.

References

Aitken, C.; Taroni, F.; and Garbolino, P. 2003. A graphical model for the evaluation of cross-transfer evidence in DNA profiles. *Theoretical Population Biology* 63(3):179–190.

Bobrow, D.; Falkenhainer, B.; Farquhar, A.; Fikes, R.; Forbus, K.; Gruber, T.; Iwasaki, Y.; and Kuipers, B. 1996. A compositional modeling language. In *Proceedings of the 10th International Workshop on Qualitative Reasoning about Physical Systems*, 12–21.

Coghill, G., and Shen, Q. 2001. On the specification of multiple models for diagnosis of dynamic systems. *AI Communications* 14(2):93–104.

Console, L., and Torasso, P. 1990. Hypothetical reasoning in causal models. *International Journal of Intelligent Systems* 5(1):83–124.

Cook, R.; Evett, I.; Jackson, G.; Jones, P.; and Lambert, J. 1999. Case pre-assessment and review in a two-wat transfer case. *Science and Justice* 39:103–111.

de Kleer, J., and Williams, B. 1987. Diagnosing multiple faults. *Artifi cial Intelligence* 32:97–130.

Evett, I.; Jackson, G.; Lambert, J.; and McCrossan, S. 2000. The impact of the principles of evidence interpretation on the structure and content of statements. *Science and Justice* 40:233–239.

Falkenhainer, B., and Forbus, K. 1991. Compositional modeling: finding the right model for the job. *Artifi cial Intelligence* 51:95–143.

Forbus, K. 1984. Qualitative process theory. *Artifi cial Intelligence* 24:85–168.

Halliwell, J., and Shen, Q. 2002. Towards a linguistic probability theory. In *Proceedings of the 11th IEEE International Conference on Fuzzy Systems*, 596–601.

Heller, U., and Struss, P. 1998. Diagnosis and therapy recognition for ecosystems - usage of model-based diagnosis techniques. In *Proceedings of the 12th International Symposium on Computer Science for Environment Protection*.

Heller, U., and Struss, P. 2001. Transformation of qualitative dynamic models - application in hydro-ecology. In Hotz, L.; Struss, P.; and Guckenbienl, T., eds., *Intelligent Diagnosis in Industrial Applications*. 95–106.

Keppens, J., and Shen, Q. 2001. On compositional modelling. *Knowledge Engineering Review* 16(2):157–200.

Keppens, J., and Shen, Q. 2004. Compositional model repositories via dynamic constraint satisfaction with order-of-magnitude preferences. *Journal of Artifi cial Intelligence Research* 21:499–550.

Keppens, J., and Zeleznikow, J. 2003. A model based reasoning approach for generating plausible crime scenarios from evidence. In *Proceedings of the 9th International Conference on Artifi cial Intelligence and Law*, 51–59.

Levy, A.; Iwasaki, Y.; and Fikes, R. 1997. Automated model selection for simulation based on relevance reasoning. *Artifi cial Intelligence* 96:351–394.

Mortera, J.; Dawid, A.; and Lauritzen, S. 2003. Probabilistic expert systems for dna mixture profiling. *Theoretical Population Biology* 63(3):191–205.

Nayak, P., and Joskowicz, L. 1996. Efficient compositional modeling for generating causal explanations. *Artifi cial Intelligence* 83:193–227.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan-Kaufmann.

Poole, D. 1993. Probabilistic Horn abduction and Bayesian networks. *Artifi cial Intelligence* 64(1):81–129.

Rickel, J., and Porter, B. 1997. Automated modeling of complex systems to answer prediction questions. *Artifi cial Intelligence* 93:201–260.

Struss, P., and Price, C. 2003. Model-based systems in the automotive industry. *AI magazine* 24(4):17–34.