Compositional Bayesian Modelling and Its Application to Decision Support in Crime Investigation

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

Abstract

Despite increasing interest in the development of intelligent techniques to aid in the prevention and detection of crime, an important issue that has not yet been addressed by existing work is the use of knowledge based systems (KBS) to aid in the synthesis and analysis of hypothetical scenarios in major crime investigation. The main limitation of conventional KBS approaches is their lack of robustness to deal with the substantial variability of crime scenarios. This paper introduces a method to apply model based reasoning techniques to this problem. In particular, the existing compositional Bayesian modelling approach is extended and adapted to create hypothetical crime scenarios. Also, methods developed in the area of model-based diagnosis are used to support the analysis of synthesised crime scenarios.

1 Introduction

In the literature on major crime investigation and evaluation of evidence, a consensus is arising that a sound methodology should at least address the following two aspects [Jamieson, 2004; Cook *et al.*, 1998]. First, the investigator must consider all the events and situations that could have resulted in each individual piece of evidence. These can eventually be combined into hypothetical scenarios that explain the case under consideration. Second, the investigator must carefully choose those investigative actions (e.g. what questions to ask and what forensic examinations to complete) that are expected to produce the most informative evidence possible. Consequently, major crime investigation constitutes a difficult craft that is hard to master without decades of in-depth experience.

This paper examines the possibility of employing knowledge based systems to support less experienced crime investigators. This presents two important challenges: the decision support system must be able (1) to cope with the enormous variability of plausible crime scenarios; and (2) to measure the information value that investigative actions are expected to yield. In order to tackle these challenges, this paper shows a compositional modelling approach to synthesise and efficiently store a space of plausible scenarios within a Bayesian Network (BN) (a refinement of the work previously given in [Keppens and Shen, 2004a]). Furthermore, it presents an application of the maximum entropy reduction technique to determine which investigative actions are most likely to reduce doubt.

The work presented herein provides a novel approach and application domain for compositional modelling. It also complements recent developments on Bayesian model based diagnosis (BMBD) [Lucas, 2001]. While the latter work constitutes a probabilistic abductive reasoning method for diagnosis, it relies on a static Bayesian network for causally explaining observations. By generating potentially different networks with a compositional modeller, more generally applicable BMBD systems can be supported.

The remainder of the paper is organised as follows: Section 2 presents the compositional Bayesian modelling approach, Section 3 describes the application of the entropy reduction technique and Section 4 concludes the paper.

2 Scenario Space Synthesis

This section describes the synthesis method for creating Bayesian scenario spaces, with Section 2.1 addressing the knowledge representation and Section 2.2 the synthesis algorithm.

2.1 Knowledge Representation

Variables

Bayesian Networks (BNs) are an efficient and comprehensible means of describing the joint probability distribution over many variables over their respective domains. The variables are created and assigned a meaning by the compositional modeller, and their probability distributions are calculated from the combined response of the influences that affect them. To facilitate the integration of these features, the subject of the reasoner proposed herein are tuples $\langle p, D_p, v_p, \oplus \rangle$. Each such tuple corresponds to a variable, which is identified by a *predicate* p, which has a *domain* D_p of values, including a default value $v_p \in D_p$, that can be assigned to the variable, and which is associated with a *combination operator* $\oplus : D_p \times D_p \mapsto D_p$ that describes how the effects of different influences acting upon the variable are combined. For example, the tuple ⟨`hanged-self(johndoe)`, {`never`, `veryfew`, `several`}, `never`, $\max$⟩[1] corresponds to a variable that describes how often `johndoe` hanged himself prior to his death.

Most variable assignments, i.e. assignments of a value to a variable from the variable's domain, correspond to plausible states and events that are part of one or more possible scenarios. There are also special types of variable assignment that convey additional information that may aid in decision support.

[1] Here, the $\max$ operator returns the highest value of the domain assuming that the ordering of values is as follows: `never` $<$ `veryfew` $<$ `several`.

These concepts have been adapted from earlier work on abductive reasoning [Poole, 1989] and model based diagnosis [Hamscher *et al.*, 1992]. In particular, some variable assignments correspond to *evidence*. These are pieces of known information that are considered to be observable consequences of a possible crime[2]. *Facts* are pieces of known information that do no require an explanation. In practice, it is often convenient to accept some information at face value without elaborating possible justifications. For instance, when a person is charged with analysing the handwriting on the aforementioned suicide note, the status of that person as a handwriting expert is normally deemed to be a fact. *Hypotheses* are possible answers to questions that must be addressed (by the investigators), reflecting certain important properties of a scenario. Typical examples of such hypotheses include the categorisation of a suspicious death into homicidal, suicidal, accidental or natural.

Also, *assumptions* are uncertain pieces of information that can be presumed to be true for the purpose of performing hypothetical reasoning. This work considers three types of assumption: (i) *Investigative actions* are assumptions that correspond to evidence collection efforts made by the investigators. For example, a variable assignment associated with the comparison of the handwriting on a suicide note and an identified sample of handwriting of the victim is an investigative action. Note that each investigative action a is associated with an exhaustive set E_a of mutually exclusive pieces of evidence that covers all possible outcomes of a. (ii) *Default assumptions* are assumptions that are presumed true unless they are contradicted. Such assumptions are typically employed to represent the conditions that an expert produces evaluations based upon sound methodology and understanding of his/her field. (iii) *Conjectures* correspond to uncertain states and events that need not be described as consequences of other states and events.

Knowledge Base

The model fragments, herein called scenario fragments in order to convey their role as parts of (crime) scenarios, describe causal influences among types of state and event in crime scenarios. However, the consequence of any influence is not presumed to be deterministic, but governed by predefined probability distributions. Thus, the notion of scenario fragment incorporates a set of probability distributions, one for each combination of the antecedent and assumption variables. As such, scenario fragments are represented by:

$$
\begin{array}{l}
\texttt{if}\ \{p_1,\ldots,p_k\} \\
\texttt{assuming}\ \{p_l,\ldots,p_m\} \\
\texttt{then}\ \{p_n\} \\
\texttt{distribution}\ p_n\ \{ \\
\qquad\vdots \\
\quad v_1,\ldots,v_k,v_l,\ldots,v_m \texttt{->} v_{n1}:q_1,\ldots,v_{nj_n}:q_{j_n} \\
\qquad\vdots \qquad\qquad \}
\end{array}
$$

[2]Note that as evidence is herein defined as "information", it does not equal the "exhibits" presented in court. Thus, for example, a suicide note is not considered to be a piece of evidence in itself, but the conclusions of a handwriting expert who has analysed the note are.

where $\{p_1,\ldots,p_k\}$ is the set of antecedent predicates, $\{p_l,\ldots,p_m\}$ is the set of assumption predicates, p_n is the consequent predicate, each v_i is a value taken from the domain D_{p_i} of the variable identified by p_i and each q_j is a real value in the range $[0,1]$. The `if`, `assuming` and `then` components of scenario fragments respectively describe the types of antecedents, assumptions and consequent of a causal relation. The domain value that a consequent has is influenced by the values of the antecedent and assumption variables and the probability distributions defined in the `distribution` component of the fragment. In particular, each line

$$v_1,\ldots,v_k,v_l,\ldots,v_m \texttt{->} v_{n1}:q_1,\ldots,v_{nj_n}:q_{j_n}$$

defines a discrete probability distribution

$$
\begin{array}{l}
f_{p_1:v_1,\ldots,p_k:v_k,p_l:v_l,\ldots,p_m:v_m\rightarrow p_n}: \\
D_{p_n}\mapsto[0,1]: f_{v_1,\ldots,v_k,v_l,\ldots,v_m}(v_{ni})=q_{ni}
\end{array}
$$

with $i=1,\ldots,j_n$. Note that it is not required that a probability distribution is defined for each combination of values assigned to the antecedent and assumption variables in a scenario fragment. Instead, a probability distribution in which the default value of the consequent variable has a probability 1 is presumed. Here, the default probability distribution for those combinations of assignments $p_1:v_1,\ldots,p_k:v_k,p_l:v_l,\ldots,p_m:v_m$ for which no probability distribution is defined, is

$$f_{p_1:v_1,\ldots,p_k:v_k,p_l:v_l,\ldots,p_m:v_m\rightarrow p_n}(v)=\begin{cases}1 & v=v_{p_n}\\ 0 & \text{otherwise}\end{cases}$$

Thus, the following scenario states that if a victim V has petechiae on his eyes and the investigators examine V's eyes, then evidence of petechiae is discovered with a certain probability:

```
if {petechiae(eyes(V))}
assuming {examination(eyes(V))}
then {evidence(petechiae(V))}
distribution evidence(petechiae(V)) {
    true, true -> true:0.99, false:0.01}
```

In the knowledge base, inconsistencies refer to inconsistent combinations of variable assignments. As such, an inconsistency denoting that $p_1:v_1\wedge\ldots\wedge p_k:v_k$ is inconsistent is represented as:

$$\texttt{inconsistent}\ \{p_1:v_1,\ldots,p_k:v_k\}$$

For example, the following inconsistency states that a person can not be both killed by another person and by him/herself:

```
inconsistent {commits-suicide-by(V,M):true,
              commits-homicide-by(P,V,M):true}
```

Inconsistencies are treated as a special type of scenario fragment of the form:

$$
\begin{array}{l}
\texttt{if}\ \{p_1,\ldots,p_k\} \\
\texttt{then}\ \{\texttt{nogood}\} \\
\texttt{distribution nogood}\ \{ \\
\quad v_1,\ldots,v_k \texttt{->} \top:1,\ldots,\bot:0\}
\end{array}
$$

where `nogood` refers to a special type of boolean variable, that remains hidden from the user and the knowledge engineer, and is known to be false. According to this definition, any situation where $p_1 : v_1, \ldots, p_k : v_k$ requires `nogood` to be true. Consequently, the probability of $p_1 : v_1, \ldots, p_k : v_k$ given that `nogood` is false, $P(p_1 : v_1, \ldots, p_k : v_k \mid \texttt{nogood} : \bot) = 0$. And in this way, the inconsistency $p_1 : v_1, \ldots, p_k : v_k$ is modelled as an impossibility.

For example, the aforementioned inconsistency is treated as a scenario fragment of the form:

```
if { commits-suicide-by(V,M),
     commits-homicide-by(P,V,M) }
then { nogood }
distribution nogood {
     true, true -> true:1, false:0 }
```

In addition to scenario fragments, the knowledge base also contains prior distributions for assumed states and events. Prior distributions are represented by

$$\texttt{define prior}\ p\ \{v_1 : q_1, \ldots, v_j : q_j\}$$

where $\{v_1, \ldots, v_j\}$ is the domain D_p of p and $q_1, \ldots, q_j$ define a function $f_p : D_p \mapsto [0,1] : f_p(v_i) = q_i$ that is a probability distribution.

For example, the definition

```
define prior suicidal(V) {true:0.02, false:0.98}
```

specifies the prior probability distribution of a variable identified by `suicidal(V)` with the domain {`true`, `false`} and

$$f_{\texttt{suicidal(V)}}(\text{true}) = 0.02$$
$$f_{\texttt{suicidal(V)}}(\text{false}) = 0.98$$

Unless specified otherwise, a variable assignment represents an uncertain state or event. However, certain variable assignments can be associated with other types of information, such as hypotheses and evidence, in the knowledge base. Predicates identifying variables whose assignments correspond to hypotheses, facts, evidence, investigative actions and default assumptions are defined by purpose built constructs that associate certain types of predicate with one of these types of information (and corresponding to evidence sets, in the case of investigative actions). Conjectures contained in the knowledge base are identified in the `assuming` clause of the scenario fragments.

Presumptions

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

1. *Any two probability distributions taken from two scenario fragments involving the same consequent variable are independent*. Intuitively, this assumption indicates that the outcome of an influence implied by one scenario fragment is not affected by that of another.
2. *There are no cycles in the knowledge base*. This means that there is no subset of scenario fragments in the knowledge base that allow a participant to be affected by itself. This assumption is required because BNs can not represent such information as they are inherently acyclic [Pearl, 1988].

While presumption 1 is a strong assumption, and may hence reflect a significant limitation of the present work, it is required herein to efficiently compute the combined effect of a number of scenario fragments on a single variable (see 2.2). Future work will seek to relax this assumption in order to generalise further the application of the method proposed.

Example

Appendix A presents a sample knowledge base with which the remaining discussion will be illustrated. The knowledge contained within it relates to cases where a person died from hanging. To keep the example self-contained, the scope of the knowledge base has been restricted and an imaginative reader may be able to produce plausible scenarios that are not covered by this knowledge base. It does, however, contain components of a broad range of scenarios, including those where the victim committed suicide, those where the victim was forcibly hanged by a murderer and those where the victim died accidentally whilst committing an act of autoerotic asphyxiation.

2.2 Algorithm

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

Structure

The procedure to synthesise the structure of a Bayesian scenario space involves a sequence of three stages. In the first stage, a hypergraph, which is similar to an ATMS network, is constructed for a given knowledge base $\mathbf{K}$, a set S of available evidence and a set F of known facts with the following algorithm:

Algorithm 2.1: GENERATEHYPERGRAPH($\mathbf{K}$, T, F)

```
comment: Initialisation:
N ← new set;   J ← new table;
for each p ∈ S ∪ F, N ← N ∪ {p}, A ← A ∪ {p};
comment: Backward chaining:
for each substitution(σ),
⟨if P_antecent then P_assumptions then {p_c}distribution D⟩ ∈ K,
σp_n ∈ N
  do { E ← new set;
       for each p_i ∈ P_antecedent ∪ P_assumptions
         do { if σp_i ∉ G
                then { p ← instantiate(σp_i);
                       N ← N ∪ {p};
                       if p_i ∈ P_assumptions
                         then A ← A ∪ {p};
                       E ← E ∪ {p};
                else E ← E ∪ {σp_i};
       J(σp_c) ← J(σp_c){E};
comment: Forward chaining:
for each substitution(σ),
⟨if P_antecent then P_assumptions then {p_c}distribution D⟩ ∈ K,
{σp_i | p_i ∈ P_antecedent} ⊂ N
  do { if σp_c ∉ N
         then { p' ← instantiate(σp_c);
                N ← N ∪ {p'));
                E ← new set;
                for each p_i ∈ P_antecedent
                  do E ← E ∪ {σp_i};
                for each p_i ∈ P_assumptions
                  do { if σp_i ∉ N
                         then { p ← instantiate(σp_i);
                                N ← N ∪ {p}; A ← A ∪ {p}
                                E ← E ∪ {p};
                         else E ← E ∪ {σp_i};
                J(p') ← J(p') ∪ {E};
```

The *initialisation* part creates two data structures to store the generated inferences temporarily. N is a set that will contain all the variables in the DAG and A is a set that will contain all the assumptions and facts. These sets are initialised with the given evidence and facts. J is a table that denotes a collection of sets with each set $J(p)$ itself containing sets of antecedents and assumptions that justify a variable identified by p, with the cardinality of the collection being the number of variables considered.

The *backward chaining* part generates all possible explanations of the available symptoms, i.e. observed evidence, by instantiating scenario fragments in $\mathbf{K}$. For each scenario fragment, whose consequent variable matches a node in N, the predicates describing the antecedent and assumption variables are instantiated and added to N if these instances do not already exist. A set E containing the antecedent and assumption instances is added to the set of justifications associated with the consequent. Note that the matching of those predicates specifying variables in N with the predicates in the scenario fragment is accomplished with a set of substitutions σ. For the example, the node corresponding to `commits-homicide-by(_1,johndoe,hanging)` matches the predicate identifying the consequent variable of the scenario fragment

```
assuming {is-killer(P,V), chooses-homicide-method(P,M)}
then {commits-homicide-by(P,V,M)}
distribution commits-homicide-by(P,V,M) {
    true -> true:1, false:0}
```

with substitution $\sigma = \{\texttt{P}\backslash\texttt{\_1}, \texttt{V}\backslash\texttt{johndoe}, \texttt{M}\backslash\texttt{hanging}\}$. Therefore, the backward chaining phase will add two new predicates `is-killer(_1,johndoe)` and `chooses-homicide-method(_1,hanging)` to N and a justification {`is-killer(_1,johndoe)`, `chooses-homicide-method(_1,hanging)`} to the set J(`commits-homicide-by(_1,johndoe,hanging)`).

The *forward chaining* part generates all possible consequences of the explanations created in the previous phase. For each scenario fragment, whose antecedent variables match instances in N, the predicates describing the assumption and consequent variables are instantiated and added to N if they do not already exist. As in the previous phase, a set E containing the antecedent and assumption instances is added to the set of justifications associated with the consequent. For the example, based on the predicate `petechiae(eyes(johndoe))` in N and the scenario fragment

```
if {petechiae(eyes(V))}
assuming {examination(eyes(V))}
then {evidence(petechiae(V))}
distribution evidence(petechiae(V)) {
    true, true -> true:0.99, false:0.01}
```

new predicates `examination(eyes(johndoe))` and `evidence(petechiae(V)` are added to N and the justification {`examination(eyes(johndoe))`, `petechiae(eyes(johndoe))`} is added to J(`evidence(petechiae(johndoe))`). In forward chaining, inconsistencies are processed in the same manner as scenario fragments, but they all have the same nogood node as their consequent, which will be identified by $n_\perp$ in what follows.

To illustrate the application of these ideas, consider a case where a person identified as `johndoe` is found dead hanging from a rope, and let this piece of evidence be identified by the variable assignment `evidence(hanging(body(johndoe)))`:`true`. By applying Algorithm 2.1 to this given piece of evidence and the knowledge base shown in Appendix A, the hypergraph shown in Figure 1 is produced.

In the second stage, spurious nodes and justifications are removed from the hypergraph. During the backward chaining phase, sets of minimal sufficient causal justifications are generated incrementally to form plausible explanations for the available evidence. Starting from the individual pieces of evidence, conjunctions of states and events justifying the pieces of evidence are created by instantiating scenario fragments, and these states and events are in turn justified by instantiating certain other scenario fragments, and so forth. Ultimately, the states and events in the justification must themselves be justified by assumptions and/or facts. As explained in Section 2.1, assumptions and facts are the only types of information that require no further explanation. Their role as so-called root nodes in the scenario space is extended in the Bayesian scenario space as they represent the only types of information which is associated with a prior distribution. In particular, assumptions have a prior distribution as defined in the knowledge base and the prior distribution of a fact corresponding to a variable assignment $p : v$ is defined by:

$$P(p : x) = \begin{cases} 1 & \text{if } x = v \\ 0 & \text{otherwise} \end{cases}$$

where $x \in D_p$. Therefore, each root node in the hypergraph generated by this procedure must be either a fact or an assumption. However, the backward chaining phase can not guarantee this. Hence, when the procedure terminates, all those nodes, which were originally regarded as root nodes and which are not a fact or assumption and all the justifications including these nodes are deemed spurious. As the hypergraph of Figure 1 does not contain any spurious nodes, consider the following scenario fragment as a means of illustration:

```
if {victim(Victim)}
assuming {
    suspect(Perpetrator),
    fight(Perpetrator,Victim),
    fight(Victim,Perpetrator)}
then {transfer(fibres,Victim,Perpetrator)}

if {victim(Victim)}
assuming {
    suspect(Perpetrator),
    fight(Perpetrator,Victim),
    fight(Victim,Perpetrator)}
then {transfer(fibres,Perpetrator,Victim)}
```

Given evidence `transfer(fibres, _1, johndoe)`, the symbolic scenario space generator will create the following information: `victim(johndoe)`, `suspect(johndoe)`, `victim(_1)`, `suspect(_1)`, `fight(_1, johndoe)` and `fight(johndoe, _1)`. Here, `victim(johndoe)` : $\top$ is a fact. Furthermore, `suspect(johndoe)` and `suspect(_1)` correspond to assumption nodes, where `suspect(johndoe)` : $\top$ should be rendered impossible by means of an inconsistency. `victim(_1)` is neither fact nor assumption, and it is not further justified. Therefore, the node containing `victim(_1)` is spurious and must be removed.

Spurious nodes are ignored in an ATMS. However, if the belief propagation algorithm of a BN attempts to take them into account, incorrect results or errors will be produced as spurious nodes have no prior probability distribution. Therefore, all spurious nodes and justifications must be removed from the hypergraph. The following procedure recursively removes from a given hypergraph $\langle N, A, J\rangle$ all root nodes that do not correspond to a fact/assumption and the justifications in which these nodes occur. The procedure terminates when each root node in $\langle N, A, J\rangle$ corresponds to either a fact or an assumption. In effect, this procedure deletes all spurious nodes and justifications from the hypergraph.

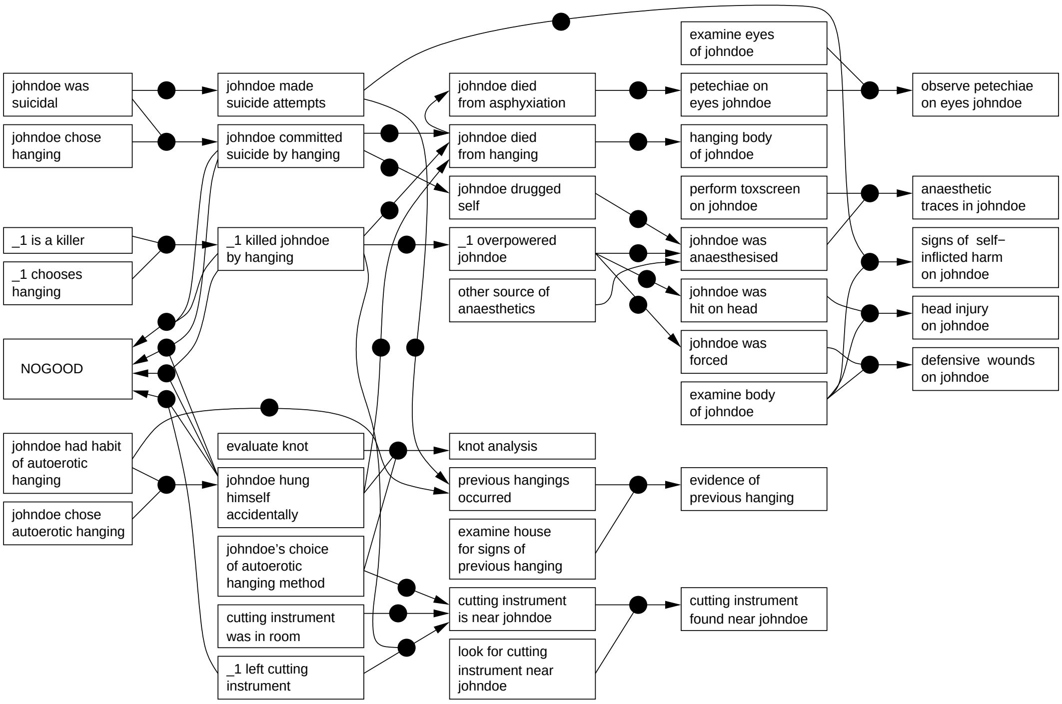

Figure 1: Sample scenario subspace structure

Algorithm 2.2: SPURIOUSNODEREMOVAL($\langle N, A, J\rangle, \mathbf{K}$)

for each $n \in N$, $[(\nexists \text{substitution}(\sigma), J(n) = \emptyset, n \notin A$
do $\begin{cases} N \leftarrow N \backslash \{n\}; \\ \textbf{for each } n', E \in J(n'), (n \in E) \wedge (n' \in N) \\ \quad \textbf{do } J(n') \leftarrow J(n') \backslash E; \end{cases}$

Finally, in the third stage, the hypergraph $\langle N, A, J\rangle$ is collapsed into a DAG by means of the following procedure:

Algorithm 2.3: CREATEDAG($\langle N, A, J\rangle$)

$G \leftarrow$ new DAG;
for each $n \in N$
do $\begin{cases} \text{add}(G, n); \\ \textbf{for each } n' \in (\bigcup_{E \in J(n)} E) \\ \quad \textbf{do } \text{add}(G, \text{arc}(n', n)); \end{cases}$

The resulting DAG forms the structure of a BN that is extended with conditional probability tables as described in the next section.

Conditional Probability Tables

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, thereby inhibiting the manual specification of the conditional probabilities. For example, the probability distribution of the amount of a particular anaesthetic in the blood of a victim's body can be affected by self-medication, consumption of a spiked drink, surgery, etc.

Using the proposed method to derive the structure of a BN from knowledge, a set $J(p_n) = \{J_1, \ldots, J_r\}$, containing sets of justifying variables is constructed for each predicate p_n, where each set of justifying variables, $J_i \in J(p_n)$, is associated an instantiated scenario fragment C_i. Each C_i contains a set of probability distributions describing how the value of the variable identified by p_n is affected by assignments to the variables in J_i. Let A be a set of value assignments to the variables in J_i or a superset thereof. Then, the probability that C_i causes p_n to take value $v \in D_{p_n}$ is denoted by $P(A \stackrel{C_i}{\rightarrow} p_n : v)$.

The set $P = \{p_1, \ldots, p_s\}$ of immediate parent variables in the generated DAG is derived by computing $J_1 \cup \ldots \cup J_r$. Let A be a set of assignments $\{p_1 : v_1, \ldots, p_s : v_s\}$, where each $v_i \in D_{p_i}$, to the parent variables of p_n in the DAG. It is clear p_n will be assigned c, with $c \in D_{p_n}$, whenever the causal influences described by the scenario fragments $C_1, \ldots, C_k$ result in a collection of outcomes $c_1, \ldots, c_k$ whose combined effect $c_1 \oplus \ldots \oplus c_k$ equals c. Thus, the probability that $p_n : c$ given A is specified by:

$$P(p_n : c \mid A) = P\left[\bigvee_{c_1 \oplus \ldots \oplus c_k = c} \left(\bigwedge_{i=1,\ldots,k} (A \stackrel{C_i}{\rightarrow} p_n : c_i)\right)\right] \quad (1)$$

According to (1), computing $P(p_n : c \mid A)$ involves calculating the likelihood of a combination of events described by a disjunctive normal form (DNF) expression. Because the occurrence of different combinations of outcomes $c_1, \ldots, c_k$ of the scenario fragments $C_1, \ldots, C_k$ involves mutually exclusive events, the calculation can be resolved by adding the probabilities of the conjuncts in (1):

$$P(p_n : c \mid A) = \sum_{c_1 \oplus \ldots \oplus c_k = c} P\left(\bigwedge_{i=1,\ldots,k} (A \stackrel{C_i}{\rightarrow} p_n : c_i)\right) \quad (2)$$

From presumption 1, the outcomes of different scenario fragments (with the same consequent), in case of a given set of assignments of the antecedent and assumption variables, correspond to independent events. Therefore, the probability of the conjunctions in (2) is equal to the product of the probabilities of their conjuncts, and (2) is calculated as follows:

$$P(p_n : c \mid A) = \sum_{c_1 \oplus \ldots \oplus c_k = c} \left(\prod_{i=1,\ldots,k} P(A \stackrel{C_i}{\rightarrow} p_n : c_i) \right) \quad (3)$$

Consider, for example, the following two scenario fragments, which are part of the probabilistic knowledge base from which its symbolic counterpart as presented in Figure 1 can be generated:

```
if { autoerotic-hanging-habit(V) }
then { previous-hanging(V) }
distribution previous-hanging(V) {
    true -> never:0.1, veryfew:0.4, several:0.5 }

if { previous-suicide-attempts(V) }
then { previous-hanging(V) }
distribution previous-hanging(V) {
    true -> never:0.7, veryfew:0.29, several:0.01 }
```

where `autoerotic-hanging-habit(V)` and `previous-suicide-attempts(V)` correspond to boolean variables, and `previous-hanging(V)` to a variable taking values from the domain $\{$`never`, `veryfew`, `several`$\}$ defined over combination operator $\max$. Then, the probabilities of assignments to `previous-hanging(V)`, given that `autoerotic-hanging-habit(V)` and `previous-suicide-attempts(V)` are assigned $\top$, can be computed as follows:

For notational convenience, let p_1, p_2 and p_3 respectively denote `autoerotic-hanging-habit(johndoe)`, `previous-suicide-attempts(johndoe)`, and `previous-hanging(johndoe)`, and let the above two scenario fragments be named C_1 and C_2. Then, the probabilities in the scenario fragments involved are assigned as:

$$P(p_1 : \top \stackrel{C_1}{\rightarrow} p_3 : \texttt{never}) = 0.1$$
$$P(p_1 : \top \stackrel{C_1}{\rightarrow} p_3 : \texttt{veryfew}) = 0.4$$
$$P(p_1 : \top \stackrel{C_1}{\rightarrow} p_3 : \texttt{several}) = 0.5$$
$$P(p_2 : \top \stackrel{C_2}{\rightarrow} p_3 : \texttt{never}) = 0.7$$
$$P(p_2 : \top \stackrel{C_2}{\rightarrow} p_3 : \texttt{veryfew}) = 0.29$$
$$P(p_2 : \top \stackrel{C_2}{\rightarrow} p_3 : \texttt{several}) = 0.01$$

According to (3), all combinations of outcomes c_1 and c_2 of scenario fragments C_1 and C_2, with $c_1, c_2 \in \{\texttt{never}, \texttt{veryfew}, \texttt{several}\}$ and `never¡veryfew¡several`, and such that $\max(c_1, c_2) =$ `veryfew`, must be considered. There are three such combinations: $\{c_1 : \texttt{veryfew}, c_2 : \texttt{veryfew}\}$, $\{c_1 : \texttt{never}, c_2 : \texttt{veryfew}\}$ and $\{c_1 : \texttt{veryfew}, c_2 : \texttt{never}\}$. Hence, $P(p_3 : \texttt{veryfew} | p_1 : \top, p_2 : \top)$ can be computed as follows:

$$\begin{aligned}
&P(p_3 : \texttt{veryfew} | p_1 : \top, p_2 : \top) \\
&= P(p_1 : \top \stackrel{C_1}{\rightarrow} p_3 : \texttt{veryfew}) \times P(p_2 : \top \stackrel{C_2}{\rightarrow} p_3 : \texttt{veryfew}) + \\
&\quad P(p_1 : \top \stackrel{C_1}{\rightarrow} p_3 : \texttt{never}) \times P(p_2 : \top \stackrel{C_2}{\rightarrow} p_3 : \texttt{veryfew}) + \\
&\quad P(p_1 : \top \stackrel{C_1}{\rightarrow} p_3 : \texttt{veryfew}) \times P(p_2 : \top \stackrel{C_2}{\rightarrow} p_3 : \texttt{never}) \\
&= 0.4 \times 0.29 + 0.1 \times 0.29 + 0.4 \times 0.7 = 0.425
\end{aligned}$$

Similarly, it can be shown that

$$P(p_3 : \texttt{never} | p_1 : \top, p_2 : \top) = 0.07$$
$$P(p_3 : \texttt{several} | p_1 : \top, p_2 : \top) = 0.505$$

3 Scenario Space Analysis

Once constructed, the Bayesian scenario space can be analysed in conjunction with the symbolic one to compute effective evidence collection strategies. The concepts of evidence, hypotheses, assumptions and facts are still employed in the Bayesian scenario space, but they now refer to variable assignments instead of predicates. For implementational simplicity, hypotheses and investigative actions are assumed to be represented by (truth) assignments to boolean variables (although this will be extended in future work).

While the likelihood ratio approach can be extended to deal with more than two hypotheses (for example by computing multiple likelihood ratios or a likelihood ratio comparing combinations of hypotheses) [Aitken and Taroni, 2004], it is not clear how these extensions can employed to compute a metric of doubt over multiple positions. The benefit of such a metric is that it enables a decision support system to order different evidence collection strategies in order of their effectiveness in reducing doubt between multiple hypotheses. An alternative approach based on information theory is proposed here.

The work will be illustrated by means of probabilities derived from a BN which has been generated by means of the the techniques of Section 2 using the knowledge base of Appendix A (and which is a Bayesian representation of the symbolic scenario space given in Figure 1).

3.1 Hypothesis sets and query types

Instead of two hypotheses, the approach aims to evaluate evidence in relation to a set H of hypotheses. This set must be exhaustive and the hypotheses within it mutually exclusive. H is *exhaustive* if one of the hypotheses in the set is guaranteed to be true, ensuring that the approach will evaluate the scenario space entirely, without ignoring any plausible scenarios. The hypotheses in a set are *mutually exclusive* if no pair of hypotheses taken from the set can be true simultaneously. This property ensures that the approach is not biased.

In this work, hypothesis sets are predefined in the knowledge base along with a precompiled taxonomy of *query types*. Query types represent important questions that the investigators need to address, such as the type of death of victim in a suspicious death case, or the killer of a victim in a homicide case. Query types are identified with a predicate describing it and they may be associated with a set of predicates identifying the hypothesis variables. For example, the following two query type definitions

```
define query type {
    unifiable = type-of-death(P),
    hypotheses = {homicidal-death(P),suicidal-death(P),
                  accidental-death(P),natural-death(P)}}

define query type {
    unifiable = killer-of(P),
    hypotheses = {killed(Q,P)}}
```

are respectively associated with the following hypothesis sets:

$$\begin{aligned} H_1 = \{&\texttt{homicidal-death(johndoe)} : \top, \\ &\texttt{suicidal-death(johndoe)} : \top, \\ &\texttt{accidental-death(johndoe)} : \top, \\ &\texttt{natural-death(johndoe)} : \top\} \\ H_2 = \{&\texttt{killed(mr-hyde,mary-kelly)} : \top, \\ &\texttt{killed(jack-the-ripper,mary-kelly)} : \top, \\ &\texttt{killed(\_1,mary-kelly)} : \top, \\ &\texttt{killed(none,mary-kelly)} : \top\} \end{aligned}$$

It is the responsibility of the knowledge engineer to ensure that the hypotheses sets generated in this way meet the exhaustiveness and mutual exclusivity criteria. These criteria can be satisfied for any given set $P = \{p_1, \ldots, p_n\}$ of predicates identifying hypotheses variables. Exhaustiveness can be assured by extending P with an additional predicate p_{n+1} and adding a probabilistic scenario fragment that enforces $p_{n+1} : \top$ with likelihood 1 if $p_1 : \bot, \ldots, p_n : \bot$, and $p_{n+1} : \bot$ with likelihood 1 otherwise:

if $\{p_1, \ldots, p_n\}$
then $\{p_{n+1}\}$
distribution p_{n+1} $\{\bot, \ldots, \bot \texttt{->} \top : 1, \bot : 0\}$

The mutual exclusivity criterion can be easily attained by adding inconsistencies for each pair of hypotheses:

inconsistent $\{p_i : \top, p_j : \top\}$

3.2 Entropy

The work here employs an information theory based approach, which is widely used in areas such as machine learning [Mitchell, 1997] and model based diagnosis [Hamscher *et al.*, 1992]. Information theory utilises a measurement of doubt over a range of choices, called entropy. Applied to the present problem, the *entropy* over an exhaustive set of mutually exclusive hypotheses $H = \{h_1, \ldots, h_m\}$ is given by:

$$\epsilon(H) = -\sum_{h \in H} P(h) \log P(h)$$

where the values $P(h)$ can be computed by means of conventional BN inference techniques. Intuitively, entropy can be interpreted as lack of information. Under the exhaustiveness and mutual exclusivity conditions, it can be shown that $\epsilon(H)$ reaches its highest value (which corresponds to a total lack of information) when $P(h_1) = \ldots = P(h_m) = \frac{1}{m}$ and $\epsilon(H)$ reaches 0 (which corresponds to a totally certain situation) when all $P(h_i)$, with $i = 1, \ldots, m$, equal 0 or 1.

In crime investigation, additional information is created through evidence collection. Thus, the entropy metric of interest for the purpose of generating evidence collection strategies is the entropy over a set of hypotheses H, given a set $E = \{e_1 : v_1, \ldots, e_n : v_n\}$ of pieces of evidence:

$$\epsilon(H \mid E) = -\sum_{h \in H} P(h \mid E) \log P(h \mid E) \qquad (4)$$

where the values $P(h \mid E)$ can, again, be computed by means of conventional BN inference techniques. For the example problem from the sample scenario space, the following probabilities can be computed, with E_1 containing hanging-dead-body(johndoe) : $\top$ and nogood : $\bot$:

$$\begin{aligned} P(\texttt{homicidal-death(johndoe)} \mid E_1) &= 0.22 \\ P(\texttt{suicidal-death(johndoe)} \mid E_1) &= 0.33 \\ P(\texttt{accidental-death(johndoe)} \mid E_1) &= 0.45 \end{aligned}$$

Thus, as an instance,

$$P(H_1 \mid E_1) = -(0.22 \log 0.22 + 0.33 \log 0.33 + 0.45 \log 0.45) = 0.46$$

A useful evidence collection strategy involves selecting investigative actions from a given set A according to the following criterion:

$$\min_{a \in A} E(\epsilon(H \mid E), a) \qquad (5)$$

Note that the entropy values calculated by equation (4) are affected by the prior distributions assigned to assumptions, as described in 2.1. Within the context of evidence evaluation (which is the conventional application of the likelihood ratio approach), this is a controversial issue as decisions regarding the likelihood of priors, such as the probability that a victim had autoerotic hanging habits, are a matter for the courts to decide on. In the context of an investigation, however, these prior distributions may provide helpful information often ignored by less experienced investigators. For example, the probability of suicides or autoerotic deaths are often underestimated. As such, decision criterion (5) is a useful means of deciding on what evidence to collect next. Yet, the minimal entropy decision rule does not yield information that should be used for evidence evaluation in court.

3.3 Minimal entropy-based evidence collection

Let a denote an investigative action and E_a be a set of the variable assignments corresponding to different possible outcomes of a (i.e. the pieces of evidence that may result from the investigative action). The expected posterior entropy (EPE) after performing a can then be computed by calculating the average of the posterior entropies under different outcomes $e \in E_a$, weighted by the likelihood of obtaining each outcome e (given the available evidence):

$$E(\epsilon(H \mid E), a) = \sum_{e \in E_a} P(e \mid a : \top, E) \epsilon(H \mid E \cup \{a : \top, e\}) \qquad (6)$$

The ongoing example contains an investigative action a =test-toxicology(johndoe): $\top$, representing a toxicology test of johndoe searching for traces of anaesthetics and a corresponding set of outcomes $E_a = \{\texttt{toxscreen(johndoe)} : \top, \texttt{toxscreen (johndoe)} : \bot\}$, respectively denoting a positive toxscreen and a negative one. Let E_2 be a set containing hanging-dead-body (johndoe) : $\top$, text-toxicology(johndoe) : $\top$ and nogood : $\bot$. Then, through exploiting the Bayesian scenario space the following can be computed:

$$
\begin{aligned}
&P(\texttt{toxscreen(johndoe)} : \top \mid E_2) = 0.17\\
&P(\texttt{toxscreen(johndoe)} : \bot \mid E_2) = 0.83\\
&P(\texttt{homicidal-death(johndoe)} \mid\\
&E_2 \cup \{\texttt{toxscreen(johndoe)} : \top\}) = 0.40\\
&P(\texttt{suicidal-death(johndoe)} \mid\\
&E_2 \cup \{\texttt{toxscreen(johndoe)} : \top\}) = 0.49\\
&P(\texttt{accidental-death(johndoe)} \mid\\
&E_2 \cup \{\texttt{toxscreen(johndoe)} : \top\}) = 0.11\\
&P(\texttt{homicidal-death(johndoe)} \mid\\
&E_2 \cup \{\texttt{toxscreen(johndoe)} : \bot\}) = 0.19\\
&P(\texttt{suicidal-death(johndoe)} \mid\\
&E_2 \cup \{\texttt{toxscreen(johndoe)} : \bot\}) = 0.44\\
&P(\texttt{accidental-death(johndoe)} \mid\\
&E_2 \cup \{\texttt{toxscreen(johndoe)} : \bot\}) = 0.38
\end{aligned}
$$

Intuitively, these probabilities can be explained as follows. In a homicide situation, anaesthetics may have been used by the murderer to gain control over johndoe, and in a suicide case, johndoe may have used anaesthetics as part of the suicide process. In the accidental (autoerotic) death case, there is no particular reason for johndoe to be anaesthetised. Therefore, the discovery of traces of anaesthetics in johndoe's body supports both the homicidal and suicidal death hypotheses whilst disaffirming the accidental death hypothesis. By means of these probabilities, the EPEs can be computed as the following instance:

$$E(\epsilon(H \mid E_1), a) = 0.17 \times 0.41 + 0.83 \times 0.45 = 0.45$$

The investigative action that is expected to provide the most information is the one that minimises the corresponding EPE. For example, Table 1 shows a number of possible investigative actions that can be undertaken (in column 1) and the corresponding EPEs in the sample Bayesian scenario space (in column 2) computed on the assumption that the aforementioned toxicology screen yielded a positive result. The most effective investigative actions in this case are a knot analysis and an examination of the body. This result can be intuitively explained by the fact that these investigative actions are effective at differentiating between homicidal and suicidal deaths, the most likely hypotheses if anaesthetics have been discovered in the body.

3.4 Extensions

While the approach presented above is itself a useful extension of the likelihood ratio approach, several further improvements are proposed.

Local optima and action sequences

Although the minimum EPE evidence collection technique guarantees to return an effective investigative action, it does not ensure globally optimal evidence collection. This limitation is inherent to any one step lookahead optimisation approach. The likelihood of obtaining poor quality locally optimal evidence collection strategies can be reduced by considering the EPEs after performing a sequence of actions $a_1, \ldots, a_v$ (of course, with incurred overheads over computation):

$$
\begin{aligned}
&E(\epsilon(H \mid E), a_1, \ldots, a_v)\\
&= \sum_{e_1 \in E_{a_1}} \cdots \sum_{e_v \in E_{a_v}} P(e_1, \ldots, e_v \mid a_1 : \top, \ldots, a_v : \top, E) \quad (7)\\
&\qquad \epsilon(H \mid e_1, a_1 : \top, \ldots, e_v, a_v : \top, E)
\end{aligned}
$$

In order to determine $E(\epsilon(H \mid E), a_1, \ldots, a_v)$, equation (7) can be simplified as follows:

$$
\begin{aligned}
&E(\epsilon(H \mid E), a_1, \ldots, a_v)\\
&= \sum_{e_1 \in E_{a_1}} \cdots \sum_{e_v \in E_{a_v}} \frac{P(e_1, \ldots, e_v, a_1 : \top, \ldots, a_v : \top, E)}{a_1 : \top, \ldots, a_v : \top, E}\\
&\qquad \epsilon(H \mid E \cup \{e_1, a_1 : \top, \ldots, e_v, a_v : \top\})\\
&= \sum_{e_1 \in E_{a_1}} \cdots \sum_{e_v \in E_{a_v}} (\prod_{i=1}^{v} \frac{P(e_i, a_1 : \top, \ldots, a_v : \top, E)}{a_1 : \top, \ldots, a_v : \top, E})\\
&\qquad \epsilon(H \mid E \cup \{e_1, a_1 : \top, \ldots, e_v, a_v : \top\})\\
&= \sum_{e_1 \in E_{a_1}} \cdots \sum_{e_v \in E_{a_v}} (\prod_{i=1}^{v} P(e_i \mid a_1 : \top, \ldots, a_v : \top, E))\\
&\qquad \epsilon(H \mid E \cup \{e_1, a_1 : \top, \ldots, e_v, a_v : \top\})
\end{aligned}
$$

Multiple evidence sets

Certain investigative actions may be associated with multiple sets of evidence. For example, a careful examination of the body of a man found hanging may yield various observations such as petechiae on the eyes, defensive wounds on the hands and lower arms and various types of discolouration of the body. The consequences of some types of investigative action, e.g. the examination of a dead body, are better modelled by multiple evidence sets since the resulting symptoms may occur in any combination of such pieces of evidence. The above approach can be readily extended to account for this by computing the EPEs after performing action a with associated evidence sets $E_{a,1}, \ldots, E_{a,w}$:

$$
\begin{aligned}
&E(\epsilon(H \mid E), a)\\
&= \sum_{e_1 \in E_{a,1}} \cdots \sum_{e_w \in E_{a,w}} P(e_1, \ldots, e_w \mid a : \top, E)\\
&\qquad \epsilon(H \mid e_1, \ldots, e_w, a : \top, E)\\
&= \sum_{e_1 \in E_{a,1}} \cdots \sum_{e_w \in E_{a,w}} (\prod_{i=1}^{w} P(e_i \mid a : \top, E))\\
&\qquad \epsilon(H \mid E \cup \{e_1, \ldots, e_w, a : \top\})
\end{aligned}
$$

Investigative action	EPE	NEER	REER
Knot analysis	0.30	26%	29%
Examine body	0.33	17%	19%
Search for cutting instrument	0.36	13%	14%
Search for signs of previous hangings	0.41	1.3%	1.5%
Check eyes for petechiae	0.46	0%	0%

Table 1: Evaluation of investigative actions

Multiple hypothesis sets
Finally, it may also be useful to consider multiple hypothesis sets instead of just one. This enables the decision support system (DSS) to propose evidence collection strategies that are effective at answering multiple queries. To consider multiple hypothesis sets $H_1, \ldots, H_t$ by measuring entropy over these sets, given a set of pieces of evidence E:

$$\begin{aligned}
&\epsilon(H_1, \ldots, H_t \mid E) \\
&= -\sum_{h_1 \in H_1} \cdots \sum_{h_t \in H_t} P(h_1, \ldots, h_t \mid E) \log P(h_1, \ldots, h_t \mid E) \\
&= -\sum_{h_1 \in H_1} \cdots \sum_{h_t \in H_t} \big(\prod_{i=1}^{t} P(h_i \mid E)\big) \log\big(\prod_{i=1}^{t} P(h_i \mid E)\big)
\end{aligned}$$

3.5 User interface

While a detailed discussion of the user interface developed for the present DSS system is beyond the scope of this paper, it is important to point out that a mere representation of the outcomes of the decision rules is inadequate for the objectives of the DSS. Investigators may have a number of considerations that are beyond the scope of the current DSS. These include perishability of evidence, legal restrictions, limitations on resources and overall workload. Therefore, the DSS is devised to list alternative evidence collection strategies in increasing order of EPEs.

The benefits of each strategy is indicated by either the *normalised expected entropy reduction* (NEER) or the *relative expected entropy reduction* (REER). The NEER represents the reduction in EPE, as a consequence of performing an investigative action a (i.e. $\epsilon(H \mid E) - E(\epsilon(H \mid E), a)$) as a proportion of the maximal entropy under total lack of information, and as such, it provides a means of assessing case progress:

$$NEER(H \mid E, a) = \frac{\epsilon(H \mid E) - E(\epsilon(H \mid E), a)}{\epsilon(H)}$$

The REER represents EPE reduction as a proportion of the entropy under the current set of available evidence, and as such, it focuses on the relative benefits of each alternative investigative action possible:

$$REER(H \mid E, a) = \frac{\epsilon(H \mid E) - E(\epsilon(H \mid E), a)}{\epsilon(H \mid E)}$$

These calculations are illustrated in Table 1 for the running example. As mentioned previously, this table presents the evaluation of a number of investigative actions after traces of anaesthetics have been discovered in `johndoe`'s body. The second column of this table displays the EPEs for investigative action while the third and fourth columns show the corresponding NEER and REER values respectively.

4 Conclusions and Future Work

This paper has presented a novel application of model based reasoning techniques to crime investigation. By means of a refinement of a compositional modelling approach for generating Bayesian Networks (BNs), the work allows for the creation of a BN that represents a space of plausible scenarios which can explain the available evidence. This use of compositional modelling enables the resulting decision support system to deal with widely varying circumstances without having to rely on an overly large knowledge base. The information contained in the BN is exploited to produce suitable evidence collection strategies that are expected to yield the most valuable information regarding significant hypotheses about the case at hand. The system provides a useful tool for aiding inexperienced major crime investigators in speculating about all plausible causes of the evidence available in a case, and in devising useful strategies to continue the investigation.

While the proposed approach presented herein offers very useful functionalities for DSS, a number of further improvements are possible. As the probability distributions in the scenario fragments refer to subjective assessments by experts of the likely outcomes, which are described in terms of vague concepts, the use of numeric probabilities conveys an inappropriate degree of precision. It would be more appropriate to incorporate a measurement of imprecision within the probability distributions. A number of approaches can provide a means of representing and reasoning with such imprecision, such as second-order probability theory [de Cooman, 2002; Goodman and Nguyen, 1999; Walley, 1997] and linguistic probability theory [Halliwell *et al.*, 2003]. Investigation into the use of symbolic probabilities forms an interesting immediate future work.

Another important consideration is that typical applications of this work involve reasoning about hypothetical scenarios that occur in time and space. The likelihood of such scenarios is not only affected by the observed symptoms or evidence, but also by constraints on the time and space in which the events in the scenarios occur. Therefore, further research into incorporating temporal and spatial reasoning in this framework is of significant relevance to this work.

Other important future work concerns relaxing two important assumptions made within this work: 1) probability distributions governing the outcomes of different causal influences (and hence represented in distinct scenario fragments) that affect the same variable must be independent, and 2) the effects of all causal influences affecting the same variable must be combinable using a single composition operator. It has been argued in this paper that these issues can be overcome by adding appropriate variables to the scenario fragments in question and that the inconvenience posed by these additional variables is far outweighed by the benefits of compositionality of scenario fragments. However, the knowledge representation scheme adopted seems to allow the aforementioned assumptions to be relaxed. For example, information on the correlation between causal influences, specified by scenario fragments, could be added to the knowledge base, thereby explicitly representing how influences are interdependent. Yet, exactly how this may be implemented requires considerable further studies. Also, multiple composition operators can be allowed by defining rules of composition, as in the work on compositional model repositories [Keppens and Shen, 2004b].

Acknowledgements

This work has been supported in part by UK EPSRC grant GR/S63267. We are very grateful to Colin Aitken and Burkhard Schafer for helpful discussions and assistance, whilst taking full responsibility for the views expressed herein.

References

[Aitken and Taroni, 2004] C. Aitken and F. Taroni. *Statistics and the evaluation of evidence*. Wileys, 2004.

[Cook *et al.*, 1998] R. Cook, I. Evett, G. Jackson, P. Jones, and J. Lambert. A model for case assessment and interpretation. *Science and Justice*, 38:151–156, 1998.

[de Cooman, 2002] G. de Cooman. Precision - imprecision equivalence in a broad class of imprecise hierarchical uncertainty models. *Journal of Statistical Planning and Inference*, 105(1):175–198, 2002.

[Goodman and Nguyen, 1999] I. Goodman and H. Nguyen. Probability updating using second order probabilities and conditional event algebra. *Information Sciences*, 121(3–4):295–347, 1999.

[Halliwell *et al.*, 2003] J. Halliwell, J. Keppens, and Q. Shen. Linguistic bayesian networks for reasoning with subjective probabilities in forensic statistics. In *Proceedings of the 9th International Conference on Artificial Intelligence and Law*, pages 42–50, 2003.

[Hamscher *et al.*, 1992] W. Hamscher, L. Console, and J. de Kleer, editors. *Readings in Model-Based Diagnosis*. Morgan-Kaufmann, San Mateo, 1992.

[Jamieson, 2004] A. Jamieson. A rational approach to the principles and practice of crime scene investigation: I. principles. *Science & Justice*, 44(1):3–7, 2004.

[Keppens and Shen, 2004a] J. Keppens and Q. Shen. Causality enabled compositional modelling of bayesian networks. In *Proceedings of the 18th International Workshop on Qualitative Reasoning about Physical Systems*, pages 33–40, 2004.

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

[Lucas, 2001] P. Lucas. Bayesian model-based diagnosis. *International Journal of Approximate Reasoning*, 27:99–119, 2001.

[Mitchell, 1997] T.M. Mitchell. *Machine Learning*. McGraw-Hill, 1997.

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

[Poole, 1989] D. Poole. Explanation and prediction: An architecture for default and abductive reasoning. *Computational Intelligence*, 5(2):97–110, 1989.

[Walley, 1997] P. Walley. Statistical inferences based on a second-order possibility distribution. *International Journal of General Systems*, 9:337–383, 1997.

A Sample Knowledge Base

```
assuming {suicidal(V)}
then {attempted-suicide(V)}
distribution attempted-suicide(V) {
    true -> true:0.8, false 0.2}

if {attempted-suicide(V)}
then {previous-hangings-occurred(V)}
distribution previous-hangings-occurred(V) {
    true -> never:0.7, veryfew:0.29, several:0.01}

assuming {suicidal(V), chooses-suicide-method(V,M)}
then {commits-suicide-by(V,M)}
distribution commits-suicide-by(V,M) {
    true, true -> true:1, false:0}

if {commits-suicide-by(V,hanging)}
then {drugged-self(V)}
distribution drugged-self(V) {
    true -> true:0.15, false:0.85}

if {other-cause-of-anaesthetics(V)}
then {was-anaesthetised(V)}
distribution was-anaesthetised(V) {
    true -> true:1, false:0}

if {drugged-self(V)}
then {was-anaesthetised(V)}
distribution was-anaesthetised(V) {
    true -> true:1, false:0}

if {commits-suicide-by(V,M)}
then {died-from(V,M)}
distribution died-from(V,M) {
    true -> true:1, false:0}

assuming {is-killer(P,V), chooses-homicide-method(P,M)}
then {commits-homicide-by(P,V,M)}
distribution commits-homicide-by(P,V,M) {
    true, true -> true:1, false:0}

if {commits-homicide-by(P,V,M)}
then {died-from(V,M)}
distribution died-from(V,M) {
    true -> true:1, false:0}

if {commits-homicide-by(P,V,hanging)}
then {overpowers(P,V)}
distribution overpowers(P,V) {
    true -> true:0.95, false:0.05}

if {overpowers(P,V)}
then {was-anaesthetised(V)}
distribution was-anaesthetised(V) {
    true -> true:0.3, false:0.7}

if {overpowers(P,V)}
then {head-injury(V)}
distribution head-injury(V) {
    true -> true:0.3, false:0.7}

if {other-cause-of-head-injury(V)}
then {head-injury(V)}
distribution head-injury(V)
    true -> true:0.3, false:0.7}

if {overpowers(P,V)}
then {was-forced(V)}
distribution was-forced(V) {
    true -> true:0.334, false:0.666}

assuming {autoerotic-asphyxiation-habit(V),
    fatal-autoerotic-asphyxiation-method(V,hanging)}
then {accidental-autoerotic-hanging(V)
} distribution accidental-autoerotic-hanging(V) {
    true, true -> true:1, false:0
    false, true -> true:0.05, false:0.95}

if {accidental-autoerotic-hanging(V)}
then {died-from(V,hanging)}
distribution died-from(V,hanging) {
    true -> true:1, false:0}

if {commits-homicide-by(P,V,M),
    accidental-autoerotic-hanging(V)}
assuming {choice-of-autoerotic-hanging-method(V),
    evaluate-knot(V)}
then {knot-analysis(V)}
distribution knot-analysis(V) {
    true, false, detachable, true ->
not-self-made:1, detachable:0, cutable:0, undetermined:0
    true, false, cutable,    true ->
not-self-made:1, detachable:0, cutable:0, undetermined:0
    false, true, detachable, true ->
not-self-made:0, detachable:1, cutable:0, undetermined:0
    false, true, cutable,    true ->
not-self-made:0, detachable:0, cutable:1, undetermined:0}

assuming {choice-of-autoerotic-hanging-method(V)}
then {cutting-instrument-near(V)}
distribution cutting-instrument-near(V) {
    cutting-instrument -> true:1, false:0}

assuming {cutting-instrument-was-at-crime-scene(V)}
then {cutting-instrument-near(V)}
distribution cutting-instrument-near(V) {
    true -> true:1, false:0}

if {commits-suicide-by(P,V,hanging)}
assuming {leaves-cutting-instrument-near(P,V)}
then {cutting-instrument-near(V)}
distribution cutting-instrument-near(V) {
    true -> true:1, false:0}


if {autoerotic-asphyxiation-habit(V)}
then {previous-hangings-occurred(V)}
distribution previous-hangings-occurred(V) {
    true -> never:0.1, veryfew:0.4, several:0.5}

if {died-from(V,hanging)}
then {cause-of-death(V,asphyxiation)}
distribution cause-of-death(V,asphyxiation) {
    true -> true:1, false:0}

if {died-from(V,hanging)}
then {evidence(hanging(body(V)))}
distribution evidence(hanging(body(V))) {
    true -> true:1, false:0}

if {cause-of-death(V,asphyxiation)}
then {petechiae(eyes(V))}
```

```
distribution petechiae(eyes(V)) {
    true -> true:0.99, false:0.01}

if {petechiae(eyes(V))}
assuming {examination(eyes(V))}
then {evidence(petechiae(V))}
distribution evidence(petechiae(V)) {
    true, true -> true:0.99, false:0.01}

if {was-anaesthetised(V)}
assuming {toxscreen(V)}
then {evidence(anaesthetics(V))}
distribution evidence(anaesthetics(V)) {
    true, true -> true:0.95, false:0.05
    false, true -> true:0.01, false:0.99}

if {head-injury(V)}
assuming {examine(body(V))}
then {evidence(head-injury(V))}
distribution evidence(head-injury(V)) {
    true, true -> true:1, false:0
    false, true -> true:0.1, false:0.9}

if {was-forced(V)}
assuming {examine(body(V))}
then {evidence(defensive-wounds(V))}
distribution evidence(defensive-wounds(V)) {
    true,  true -> true:1, false:0
    false, true -> true:0.1, false:0.9}

if {attempted-suicide(V)}
assuming {examine(body(V))}
then {evidence(self-harm(V))}
distribution evidence(self-harm(V)) {
    true, true -> true:0.7, false:0.3
    false, true -> true:0.1, false:0.9}

if {previous-hangings-occurred(V)}
assuming {search(home(V),previous-hangings)}
then {evidence(previous-hangings(V))}
distribution evidence(previous-hangings(V)) {
    veryfew, true -> true:0.2, false:0.8
    several, true -> true:0.95, false:0.05}

if {cutting-instrument-near(V)}
assuming {search(near(V),cutting-instrument)}
then {evidence(cutting-instrument-near(V))}
distribution evidence(cutting-instrument-near(V)) {
    true, true -> true:0.9, false:0.1}

inconsistent {commits-suicide-by(V,M):true,
              commits-homicide-by(P,V,M):true}

inconsistent {commits-suicide-by(V,M):true,
              accidental-autoerotic-hanging(V):true}

inconsistent {commits-homicide-by(P,V,hanging):true,
              accidental-autoerotic-hanging(V):true}

inconsistent {leaves-cutting-instrument-near(P,V):true,
              accidental-autoerotic-hanging(V):true}

define prior suicidal(V) {true:0.02, false:0.98}

define prior chooses-suicide-method(V,hanging) {
    true:0.1, false:0.9}

define prior is-killer(P,V) {
    true:0.01, false:0.99}

define prior chooses-homicide-method(P,hanging) {
    true:0.05, false:0.95}

define prior autoerotic-asphyxiation-habit(V) {
    true:0.025, false:0.975}

define prior fatal-autoerotic-hanging(V) {
    true:0.01, false:0.99}

define prior other-cause-of-anaesthetics(V) {
    true:0.05, false:0.95}

define prior other-cause-of-head-injury(V) {
    true:0.05, false:0.95}

define prior leaves-cutting-instrument-near(P,V) {
    true:0.5, false:0.5}
```