

Aberystwyth University

Compositional Bayesian Modelling for Computation of Evidence Collection Strategies

Keppens, Jeroen; Shen, Qiang; Price, Christopher John

Published in:
Applied Intelligence

DOI:
[10.1007/s10489-009-0208-5](https://doi.org/10.1007/s10489-009-0208-5)

Publication date:
2011

Citation for published version (APA):
Keppens, J., Shen, Q., & Price, C. J. (2011). Compositional Bayesian Modelling for Computation of Evidence Collection Strategies. *Applied Intelligence*, 35(1), 134-161. <https://doi.org/10.1007/s10489-009-0208-5>

General rights

Copyright and moral rights for the publications made accessible in the Aberystwyth Research Portal (the Institutional Repository) are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Aberystwyth Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Aberystwyth Research Portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

tel: +44 1970 62 2400
email: is@aber.ac.uk

Compositional Bayesian Modelling for Computation of Evidence Collection Strategies

Jeroen Keppens*, Qiang Shen[†] and Chris Price[‡]

Abstract

As forensic science and forensic statistics become increasingly sophisticated, and judges and juries demand more timely delivery of more convincing scientific evidence, crime investigation is becoming progressively more challenging. In particular, this development requires more effective and efficient evidence collection strategies, which are likely to produce the most conclusive information with limited available resources. Evidence collection is a difficult task, however, because it necessitates consideration of: a wide range of plausible crime scenarios, the evidence that may be produced under these hypothetical scenarios, and the investigative techniques that can recover and interpret the plausible pieces of evidence. A knowledge based system (KBS) can help crime investigators by retrieving and reasoning with such knowledge, provided that the KBS is sufficiently versatile to infer and analyse a wide range of plausible scenarios. This paper presents such a KBS. It employs a novel compositional modelling technique that is integrated into a Bayesian model based diagnostic system. These theoretical developments are illustrated by a realistic example of serious crime investigation.

1 Introduction

In the literature on major crime investigation and evaluation of evidence, the consensus is that a sound working methodology should at least involve the following two aspects [9, 29]. Firstly, for each piece of evidence, investigators should formulate conjectures, each providing a sufficient explanation for what has led to it. The set of conjectures considered for the evidence ought to be as complete as possible, and it needs to be defined independently of any broader hypotheses about the crime under investigation. Secondly, the pieces of evidence and corresponding conjectures should be combined to formulate plausible crime scenarios that explain all of the available evidence. Any conjectures that sufficiently explain multiple pieces of evidence should be identified, as well as any constraints between conjectures

*Department of Computer Science, King's College London, Strand, London WC2R 2LS, UK.

[†]Department of Computer Science, Aberystwyth University, Penglais Campus, Aberystwyth SY23 3DB, UK.

[‡]Department of Computer Science, Aberystwyth University, Penglais Campus, Aberystwyth SY23 3DB, UK.

While such an approach to crime investigation is sound in theory, it may be difficult to accomplish in practice. Human investigators have a tendency to restrict the conjectures they consider to those which seem to fit with preconceived assumptions [18]. This is particularly problematic in situations requiring complex hypothetical reasoning such as serious crime investigation [27]. Moreover, this approach demands considerable expertise regarding the plausible causes of evidence, the way hypothesised causes can be combined into scenarios, the (new) evidence that follows from the scenarios under consideration, and the investigative techniques that may help to recover such (new) evidence.

This paper examines the possibility of employing knowledge-based systems, in the form of a decision support system (DSS), to assist less experienced crime investigators, by providing them with the necessary expertise. In particular, the work presented here draws initial ideas from preliminary research as previously reported in [32, 34, 33, 35, 36], offering a substantially developed approach to aid in serious crime investigation with novel artificial intelligence techniques. This research combines compositional modelling techniques [34], the non-probabilistic mechanisms which underly a specific DSS as developed in [32, 36], the Bayesian Network synthesis approach that emerges from [33], and the evidence collection algorithms proposed in [35], within a common framework. It refines and formalises each of these elements to enable their integration into an implemented, coherent, prototypical software system.

Until now, the development of such a DSS has received little attention in the area of artificial intelligence and law. Instead, existing work has focussed on addressing other important problems such as: establishment of the validity of evidence [5], intelligent analysis of data warehouses [6, 30, 42, 60], identification of plausible crimes on the basis of evidence [16, 24, 25, 48, 54], hypothesis formulation for individual cases on the basis of evidence and evaluation of their likelihood [2, 9, 40], and generation of appropriate legal arguments in court proceedings on the basis of evidence [8, 38, 47, 55, 58]. The work introduced here provides a complementary alternative to this, in an attempt to avoid failures to consider crucial lines of inquiry (which have been identified as a prominent cause of miscarriages of justice [14]).

In particular, this research tackles two challenging problems in building DSS for crime investigation and evidence evaluation: (1) coping with the enormous variability of plausible crime scenarios; and (2) to estimating the information value that further investigating actions may be expected to possess. In dealing with these challenges, this paper presents a compositional modelling approach to synthesising and efficiently storing a space of plausible scenarios within a Bayesian Network (BN). This is then integrated into a novel type of Bayesian model based diagnostic system that employs a maximum expected entropy reduction technique to determine which investigating strategies are likely to produce the most conclusive evidence.

The remainder of this paper is organised as follows. Section 2 provides a summary on the use of BNs for evidence evaluation. Section 3 shows the broader architecture of the proposed DSS. Sections 4 and 5 jointly present the compositional Bayesian modelling approach to automatically generating a space of plausible scenarios in the form of a BN. Section 6 describes the application of the entropy reduction technique. Section 7 discusses the work in relation to other studies, as well as the challenges in applying the proposed approach in practice. Finally, section 8 concludes the paper.

2 Background

In order to produce effective evidence evaluation strategies, a method is required to evaluate the informative value of a piece of evidence. As argued in [51], a wide range of methodologies have been devised for this purpose. The present work will employ Bayesian belief propagation to evaluate evidence. This is because there is a substantial body of research on forensic statistics in which Bayesian networks (BN) are developed as probabilistic expert systems for evaluating specific types of forensic evidence [1, 53]. Therefore, this section presents a brief overview of this existing methodology.

Briefly, the method for applying the Bayesian approach to evaluating a piece of forensic evidence e [17] follows the following procedure:

1. Identify the *prosecution position* $p_{\text{prosecution}}$. This may be the case of a prosecution attorney after the investigation or a hypothesis of the forensic scientist or crime investigator.
2. Identify the *defence position* p_{defence} . This may be the case of the defence attorney, an explanation of a suspect, or a presumed “best defence”.
3. Build a model to compute the probability $P(e \mid p_{\text{prosecution}})$ of obtaining the given piece of evidence in the prosecution scenario, and another to compute the probability $P(e \mid p_{\text{defence}})$ of obtaining the given piece of evidence in the defence scenario. One approach to modelling these probabilities is to use BNs. BNs describe how the probability of the evidence of interest is affected by causes within and outside of the prosecution and defence scenarios.
4. Calculate the *likelihood ratio*:

$$LR = \frac{P(e \mid p_{\text{prosecution}})}{P(e \mid p_{\text{defence}})} \quad (1)$$

The ratio in (1) gives the value of the probability of the evidence if the prosecution’s scenario is true relative to the probability of the evidence if the defence scenario is true. The vertical bar $|$ denotes conditioning. The characteristic to the left of the bar is the event or hypothesis whose outcome is uncertain and for which a probability is wanted. The characteristic to the right of the bar are the events or hypotheses which are assumed known.

The greater LR is, the more support evidence e provides for the prosecution position. The closer LR is to 0 (and smaller than 1), the better e supports the defence position. If LR is around 1, the evidence provides little information about either position. As such, LR can be employed as a means for a forensic expert to make consistent statements in court about the implications of evidence and as a tool for investigators to decide the potential benefit of an expensive laboratory experiment prior to committing any resources.

The methodology of inferring and comparing the (body of) evidence that should be observed under conjectured (prosecution or defense) scenarios corresponds to the hypothetico-deductive method that is widely adopted in science, and which is gaining increased acceptance in serious crime investigation [29]. The specific use of precise probabilities is more




Figure 1: A simple Bayesian network

controversial, although it is adopted by major forensic laboratories, such as the Forensic Science Service of England and Wales [9]. Obviously, the approach is very useful when substantial data sets enable the analyst to calculate accurate estimates. This is the case in evaluating DNA evidence, for example [40]. Nevertheless, the approach can also be successfully applied to cases where the analyst has to rely on subjective probabilities, by performing a detailed sensitivity analysis [2] or by incorporating uncertainty concerning the probability estimates within the Bayesian model [21].

The likelihood ratio method is crucially dependent upon a means to compute the probabilities $P(e \mid p_{\text{prosecution}})$ and $P(e \mid p_{\text{defense}})$. As shown in [2, 10], Bayesian Networks (BN) are a particularly useful technique in this context. A BN is a directed acyclic graph (DAG) whose nodes correspond to random variables and whose arcs describe how the variables are dependent upon one another. Each variable can be assigned a value, such as "true" or "false", and each assignment of a value to a variable describes a particular situation of the real world (e.g. "Jones has an alibi" is "true"). The arcs have directions associated with them. Consider two nodes, A and B say, with a directed arc pointing from A to B . Node A is said to be the parent of B and B is said to be the child of A . Moreover, an arc from a node labelled H pointing towards a node labelled e indicates that e is dependent on H . The graph is acyclic in that it is not permitted to follow directed arcs and return to the starting position. Thus, a BN is a type of graphical model that captures probabilistic knowledge. The actual probabilistic knowledge is specified by probability distributions: a prior probability distribution $P(x_i) : D_{x_i} \mapsto [0, 1]$ for each root node x_i in the DAG and a conditional probability distribution $P(x_i \mid x_j, \dots, x_k) : D_{x_i} \times D_{x_j} \times \dots \times D_{x_k} \mapsto [0, 1]$ for each node x_i that has a set of (immediate) parent nodes $\{x_j, \dots, x_k\}$ (where D_x , the domain of variable x , denotes the set of all values that can be assigned to x).

Figure 1 illustrates these concepts by a sample BN that describes how the assumed event of a perpetrator killing a victim (x_1) is related to the possibility of discovering traces of blood on the victim matching the perpetrator's blood (x_5). The probabilistic knowledge represented by this BN is specified by probability distributions. There are two prior probability distributions $P(x_1)$ is read as the probability that the perpetrator has killed the victim and $P(x_4)$ read as the probability that the victim has the background they have, e.g. if a victim is a 20-year-old female student at a red-brick university, then $P(x_4)$ is the probability that a victim of rape is a 20-year-old female student at a red-brick university (in the absence of other information). This probability is a subjective one. There are three conditional probabilities $P(x_2 \mid x_1)$, $P(x_3 \mid x_2, x_1)$ and $P(x_5 \mid x_4, x_3)$ where the probabilities of the various possible combinations of outcomes are represented in tables. Table 1 shows such an example.

$P(x_2 x_1)$	x_1 : perpetrator killed victim	
x_2 : perpetrator had violent contact with the victim	True	False
True	0.9	0.3
False	0.1	0.7

Table 1: Example conditional probability table

Given the figures of Table 1, the subjective probability that the perpetrator had violent contact with victim (x_2 : true) given that the perpetrator killed the victim (x_1 : true) is 0.9.

BNs facilitate the computation of joint and conditional probabilities. This can be illustrated by calculating $P(x_5 : \text{true} \mid x_1 : \text{false})$. By definition,

$$P(x_5 : \text{true} \mid x_1 : \text{false}) = \frac{\sum_{v_2 \in D_{x_2}} \sum_{v_3 \in D_{x_3}} \sum_{v_4 \in D_{x_4}} P(x_1 : \text{false}, x_2 : v_2, x_3 : v_3, x_4 : v_4, x_5 : \text{true})}{P(x_1 : \text{false})}$$

Note that in general, this calculation requires a substantial number of joint probabilities $P(x_1 : \text{false}, x_2 : v_2, x_3 : v_3, x_4 : v_4, x_5 : \text{true})$. However, with the BN of Figure 1 and the corresponding conditional probability tables, the computation can be reduced to

$$\begin{aligned}
P(x_5 : \text{true} \mid x_1 : \text{false}) &= \sum_{v_2 \in D_{x_2}} \sum_{v_3 \in D_{x_3}} \sum_{v_4 \in D_{x_4}} P(x_3 : v_3 \mid x_2 : v_2, x_1 : \text{false}) \times \\
&\quad P(x_2 : v_2 \mid x_1 : \text{false}) \times \\
&= \sum_{v_2 \in D_{x_2}} P(x_2 : v_2 \mid x_1 : \text{false}) \times \sum_{v_3 \in D_{x_3}} P(x_3 : v_3 \mid x_2 : v_2, x_1 : \text{false}) \times \\
&\quad \sum_{v_4 \in D_{x_4}} P(x_5 : \text{true} \mid x_3 : v_3, x_4 : v_4)
\end{aligned}$$

As such, a BN can be employed both as a means to describe probabilistic knowledge and to calculate probabilities more efficiently. This increase in efficiency of the computation arises because the BN indicates which variables or factors are conditionally independent of each other. Thus the amount of blood traces found on the victim which match the perpetrator (x_5) is independent of whether the perpetrator has killed the victim or not (x_1) if the amount of blood traces transferred from the perpetrator to the victim is known (i.e. conditional on the knowledge of the amount of blood traces transferred from the perpetrator to the victim), x_3 . This conditional independence is indicated by the separation of the node for x_1 from x_5 by x_3 .




Figure 2: System Architecture

3 System Architecture

A drawback of the conventional Bayesian evidence evaluation approach is that it relies on BNs designed for evaluating specific types of evidence e with respect to specific hypotheses $p_{\text{prosecution}}$ and p_{defense} . For example, the BN of Figure 1 is restricted to evaluate only blood transfer evidence with respect to a homicide hypothesis. As such, the approach requires careful consideration and specification of the hypotheses and the corresponding BNs. This limits the generality and reusability of the approach. The present work addresses this issue by (i) developing of a novel inference mechanism that can synthesise a BN which describes a range of plausible scenario and corresponding hypotheses, and (ii) employing Bayesian model based diagnostic techniques to evaluate evidence and to create evidence collection strategies with regard to synthesised scenarios or hypotheses.

The overall architecture of the resulting DSS is shown in Figure 2. Given a new or ongoing investigation in which an initial set of evidence has been collected by the investigators, the *synthesis* component will generate a large BN that contains a range of plausible scenarios, called the *scenario space*, by means of a *knowledge base*. The scenario space describes how plausible states and events, pieces of evidence, investigative actions and other properties of crime scenarios are causally related to one another. Note that this representation is substantially more efficient than one that considers a separate representation for each considered scenario because all the information that is shared by different scenarios is stored only once.

Once the scenario space has been constructed, the *analysis* component can evaluate it to provide useful information, such as (i) descriptions of scenarios that may have caused the available evidence, (ii) useful explanations for the human investigator (e.g what evidence should be expected if a certain scenario were true), and (iii) investigative actions that are likely to yield useful information. The proposed investigative actions can, in turn, be used to collect evidence or to *speculate* about the possible outcomes of the investigative action and determine future actions.

4 Knowledge Representation

4.1 Plausible situation

As this work concerns hypothetical crime scenarios, much reasoning focusses on plausible situations (i.e. a particular condition or status of the world) and events (i.e. a development that changes a situation). In what follows, the term “situation” will be employed to refer to both states and events that change states.

Situations can be described at different levels of generality, depending on what is being modelled. In terms of relations between possible scenarios and a given specific case, the system supports the representation of situation *instances* at the information level, such as:

- “fibres were transferred from *jane doe* to *joe bloggs*” and
- “*many* fibres were transferred from *jane doe* to *joe bloggs*”

Generally speaking, situation instances refer to information about specific entities, such as *jane doe* and *joe bloggs*. At the knowledge level, which relates to the understanding of the crime investigation domain, the system also supports representation of situation *types*, such as:

- “fibres were transferred from *person P1* to *person P2*” and
- “*many* fibres were transferred from *person P1* to *person P2*”

Thus, situation types refer to certain general properties or relations of the classes of entities.

As the examples illustrate, quantities and truth values are also an important feature of situations and events. In a manner similar to the distinction between types and instances, the system may sometimes be interested in situations and events that denote specific quantities and sometimes it may not. Situations and events that leave open specific quantities or truth values involved, such as

- “fibres were transferred from *jane doe* to *joe bloggs*” and
- “fibres were transferred from *person P1* to *person P2*”

are referred to as *variables*. Situations that do include specific quantities, such as

- “*many* fibres were transferred from *jane doe* to *joe bloggs*” and
- “*many* fibres were transferred from *person P1* to *person P2*”

are referred to as (*variable*) *assignments*, the variable being the 'quantity' of the fibres transferred, where 'quantity' can take, say, one of three values 'none', 'few', or 'many'.

To facilitate the integration of these features, variables are defined by tuples $\langle p, D_p, v_p, \oplus \rangle$. The variable defined by a tuple $\langle p, D_p, v_p, \oplus \rangle$ is identified by a *predicate* p , has a *domain* D_p of values, including a *default value* $v_p \in D_p$, that can be assigned to the variable, and 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 reasons that will become clearer in the paper, the default value is set to be the neutral element of the combination operator in normal circumstances. For example, the tuple $\langle \text{previous-hangings}(\text{johndoe}), \{\text{never}, \text{veryfew}, \text{several}\}, \text{never}, \text{max} \rangle$ corresponds to a variable that describes the number of hangings johndoe survived before his death. Here, the max operator returns the greatest value of the domain, assuming that the ordering of values is as follows: $\text{never} < \text{veryfew} < \text{several}$.

Most variable assignments 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. These concepts have been adapted from earlier work on abductive reasoning [46] and model based diagnosis [23]. 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. Note that as evidence is 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 conclusion of a handwriting expert who has analysed the note is. *Facts* are pieces of known information that do not 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 a 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.

4.2 Knowledge Base

The objective of crime investigation is to determine (ideally beyond a reasonable doubt) the crucial features (such as the type of crime, the perpetrator(s), etc.) that have led to, or more precisely caused, the available evidence. Therefore, the domain knowledge that is relevant to crime investigation concerns the causal relations among the plausible situations.

4.2.1 Scenario fragments

These causal relations are stored in the knowledge base: The decision support system described in this paper employs a knowledge base consisting of generic causal relations, called scenario fragments, among situations. Each scenario fragment describes a phenomenon whereby a combination of situations leads to a new situation. In particular, each scenario fragment consists of (i) a rule representing *which* situation variable types are causally related, and (ii) a set of probability distributions that represent *how* the corresponding situation assignment types are related by the phenomenon in question.

Formally, scenario fragments are represented by constructs of the form:

$$\begin{array}{l} \text{if } \{p_1, \dots, p_k\} \\ \text{assuming } \{p_l, \dots, p_m\} \\ \text{then } \{p_n\} \\ \text{distribution } p_n \{ \\ \quad \vdots \\ \quad v_1, \dots, v_m \rightarrow v_{n1} : q_1, \dots, v_{nj_n} : q_{j_n} \\ \quad \vdots \quad \quad \quad \} \end{array}$$

where $\{p_1, \dots, p_k\}$ is the set of antecedent predicates, $\{p_l, \dots, p_m\}$ is the set of assumption predicates, p_n is the consequent predicate, each $v_i, i \in \{1, \dots, m\}$, 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 antecedent, assumption and consequent predicates identify situation variable types. The antecedents and assumptions of a rule refers to the situations that are required for the phenomenon described by the scenario fragment to take effect. The consequent describes the effect of the phenomenon. Note that the assumptions are situations that may be presumed to be present in a scenario because they refer to conjectures, default assumptions or investigative actions, and that the antecedents are themselves situations that must either be consequences of certain other scenario fragments or be factually true.

For example, the scenario fragment:

```

if {
  anesthetic-substance(A) }
assuming {
  is-addicted-to(P,A),
  has-recently-used(P,A) }
then {
  carries-traces-in-blood-of(P,A) }
distribution {
  true, true, true -> none:0, low:0.2, high: 0.8
  true, false, true -> none:0.6, low:0.3, high:0.1 }

```

denotes that the likelihood of the conjectures "P is addicted to an anesthetic substance A" and "P has recently used anesthetic substance A" affects that of the situation that P's blood carries traces of A.

A scenario fragment, which has the variable types p_1, \dots, p_m as its antecedents and assumptions and the variable type p_n as its consequent, must also define a probability distribution over the possible assignments of p_n , for each combination of assignments to p_1, \dots, p_m . These distributions describe how the assignment of p_n is affected by the phenomenon under different configurations of p_1, \dots, p_m . Each line of the form

$$v_1, \dots, v_m \rightarrow v_{n1} : q_1, \dots, v_{nj_n} : q_{j_n}$$

in the `distribution` component of a scenario fragment, defines such a probability distribution $D_{p_n} \mapsto [0, 1] : f(v_i) = q_i$, for the combination of antecedent and assumption assignments $p_1 : v_1, \dots, p_m : v_m$. In what follows, the probability (q_i) that the phenomenon described by scenario fragment S causes the consequent variable p_n to take value v_{ni} given the assignments $p_1 : v_1, \dots, p_m : v_m$ will be denoted by:

$$P(p_1 : v_1, \dots, p_m : v_m \xrightarrow{S} p_n : v_{ni}) = q_i$$

In the aforementioned example, let the domains of `anesthetic-substance(A)`, `is-addicted-to(P,A)` and `has-recently-used(P,A)` be $\{\text{false}, \text{true}\}$ and the domain of `carries-traces-in-blood-of(P,A)` be $\{\text{none}, \text{low}, \text{high}\}$. The scenario fragment is completed with specifications of probability distributions such as:

```

true, true, true -> none:0, low:0.2, high: 0.8

```

which states that the phenomenon described by the scenario fragment has a probability of 0.8 to contribute a high number of traces of A to the blood of P, a probability of 0.2 to contribute a low number of traces of A to the blood of P and a probability of 0 to contribute no traces of A to the blood of P. Or, as this will be denoted in the remainder of this paper:

$$P(p_1 : \text{true}, p_2 : \text{true}, p_3 : \text{true} \xrightarrow{C} p_4 : \text{high}) = 0.8$$

$$P(p_1 : \text{true}, p_2 : \text{true}, p_3 : \text{true} \xrightarrow{C} p_4 : \text{low}) = 0.2$$

$$P(p_1 : \text{true}, p_2 : \text{true}, p_3 : \text{true} \xrightarrow{C} p_4 : \text{none}) = 0$$

where C denotes the aforementioned example scenario fragment and

$p_1 = \text{anesthetic-substance}(A),$

$p_2 = \text{is-addicted-to}(P, A),$

$p_3 = \text{has-recently-used}(P, A)$ and

$p_4 = \text{carries-traces-in-blood-of}(P, A)$

Note that it is not required that a probability distribution be defined for each combination of values assigned to the antecedent and assumption variables in a scenario fragment S . 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 $p_1 : v_1, \dots, p_m : v_m$ for which no probability distribution is defined, is

$$P(p_1 : v_1, \dots, p_m : v_m \xrightarrow{S} p_n : v) = \begin{cases} 1 & v = v_{p_n} \text{ (i.e. the default value of } p_n) \\ 0 & \text{otherwise} \end{cases}$$

Let the default value of `carries-traces-in-blood-of`(P, A) be 'none'. Then, the aforementioned example scenario fragment also defines the following probability distribution implicitly:

$$P(p_1 : \text{true}, p_2 : \text{false}, p_3 : \text{false} \xrightarrow{C} p_4 : \text{high}) = 0$$

$$P(p_1 : \text{true}, p_2 : \text{false}, p_3 : \text{false} \xrightarrow{C} p_4 : \text{low}) = 0$$

$$P(p_1 : \text{true}, p_2 : \text{false}, p_3 : \text{false} \xrightarrow{C} p_4 : \text{none}) = 1$$

4.2.2 Inconsistencies

Some scenario fragments impose hard constraints on the possible scenarios. These scenario fragments are called *inconsistencies*. They contain a set of situation assignment types that are jointly impossible, and hence do not constitute any part of a sound scenario or a possible explanation. In the knowledge base, inconsistencies define inconsistent combinations of variable assignments. As such, an inconsistency denoting that $p_1 : v_1 \wedge \dots \wedge p_k : v_k$ is inconsistent is represented as:

$$\text{inconsistent } \{p_1 : v_1, \dots, p_k : v_k\}$$

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:

```
if {p1,...,pk}
then {nogood}
distribution nogood {
    v1,...,vk->⊤:1,...,⊥:0}
```

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 $p_1 : v_1, \dots, p_k : v_k$ requires `nogood` to be true. Consequently, the probability of $p_1 : v_1, \dots, p_k : v_k$ given that `nogood` is false, $P(p_1 : v_1, \dots, p_k : v_k \mid \text{nogood} : \perp) = 0$. And in this way, the inconsistency $p_1 : v_1, \dots, p_k : v_k$ 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 }
```

4.2.3 Prior distributions

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

```
define prior p {v1:q1,...,vj:qj}
```

where $\{v_1, \dots, v_j\}$ is the domain D_p of p and q_1, \dots, 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

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

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

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.

4.3 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, given the relevant antecedents and assumptions, the phenomena described by different scenario fragments that affect the same situation variable instance are independent from one 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 situation/event variable instance to be affected by itself. This assumption is required because BNs can not represent such information as they are inherently acyclic [45].

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 5.2). Future work will seek to relax this assumption in order to generalise further the application of the method proposed.

4.4 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

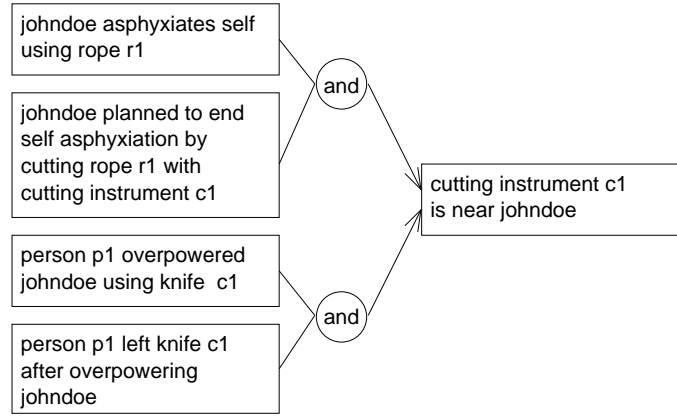


Figure 3: Partial scenario space

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.

5 Scenario Space Synthesis

A scenario space is a grouping of all the scenarios that support the available evidence, and of the evidence and hypotheses that these scenarios entail. In a typical case, there may be many scenarios that explain the given evidence and most of these may only involve minor variations from the others in the same space. Therefore, it is not sensible to describe scenarios individually. Instead, the scenario space contains the situations that constitute scenarios and how they are causally related.

Formally, a scenario space covers two sub-representations: the structural scenario space and the probabilistic scenario space:

- *The structural scenario space* denotes which combinations of situations affect other situations within the scenario space. Each situation in the structural scenario space is deemed to be either an assumption/fact or a consequent of other situations. For each sequence of situations, through the use of an assumption-based truth maintenance system [12], the structural scenario space maintains the smallest sets of other situations that can meaningfully explain it.

Consider, for example, the partial structural scenario space depicted in Figure 3. This space identifies two plausible causes for "cutting instrument c1 is near johndoe":

1. A combination of events that may be part of an accidental auto-asphyxiation scenario: "johndoe asphyxiates self using rope r1" and "johndoe planned to end self asphyxiation by cutting rope r1 with cutting instrument c1", and

2. A combination of events that may be part of a homicidal asphyxiation scenario: "person p1 overpowered johndoe using knife c1" and "person p1 left knife c1 after overpowering johndoe".


- *The probabilistic scenario space* denotes *how likely* a combination of situations affects other situations in the scenarios space. The probabilistic scenario space is a Bayesian network that refers to the same nodes as the structural scenario space. An example of a probabilistic scenario space is the Bayesian network depicted in Figure 1, with associated probability tables.

The structural and probabilistic versions of the scenario space reflect different types of information, providing complementary knowledge. While the structural scenario space does not provide a suitable means to test hypotheses or to evaluate evidence in an investigation, it enables the generation of explanatory scenarios, as illustrated in the example of Figure 3. Contrarily, while the probabilistic scenario space lacks the knowledge to produce minimally explanatory scenarios, it does, as shown in section 2, provide the necessary knowledge to test hypotheses and evaluate evidence. This section aims to demonstrate how both scenario spaces can be constructed from a given knowledge base.

5.1 Generating the structural scenario space

Informally, the scenario space structure contains:

- (i) The hypothetical causes of the available evidence. As mentioned earlier, evidence has been defined as observable consequences of scenarios. Therefore, the probability of evidence is modelled causally. This is the approach taken by most of the existing abductive diagnostic techniques [23]. The decision support system seeks to identify combinations of these causes with each combination constituting a plausible crime scenario, and to help the human investigators differentiate between the likelihood of alternative scenarios.
- (ii) Additional potential evidence, yet undiscovered, that may exist as the consequences of the aforementioned causes. The decision support system discussed here seeks to identify which investigative activities might be the most effective at differentiating between the plausibility of alternative scenarios. This is in order to discover particular types of evidence.
- (iii) The plausible alternative hypothetical causes for the secondary evidence mentioned under (ii). This is required to correctly assess the likelihood of obtaining evidence through certain investigative activities. This information is employed to help determine which investigative activities are most effective. Note that these additional hypothetical causes are not part of the diagnostic problem in the way that the hypothetical causes mentioned under (i) are. It is not until they become plausible explanations of observed evidence that they have to be considered as such. For this reason, it is not necessary to consider additional evidence potentially resulting from the causes mentioned under (iii), until these causes may have explained observed evidence.



$$\begin{aligned}
 N &= \{n_1, n_2, n_3, n_4, n_5, a_1, a_2, a_3\} \\
 A &= \{a_1, a_2, a_3\} \\
 J(n_1) &= \{\{n_2, n_3\}, \{n_4, n_5\}\} \\
 J(n_2) &= \{\{a_1\}\} \\
 J(n_3) &= \{\{a_1, a_2\}\} \\
 J(n_4) &= \{\{a_2, a_3\}\} \\
 J(n_5) &= \{\{a_3\}\}
 \end{aligned}$$

Figure 4: Sample hypergraph $\langle N, A, J \rangle$

- (iv) The combinations of causes mentioned in (i) and (iii) that are inconsistent. Certain combinations of situations and events can not be part of the same scenario because it is impossible that they are part of the same possible world. For example, when a person shoots someone across a long distance, transfer of blood splatter from the victim to the shooter as a result of the shot is impossible. Such inconsistent combinations of situations and events need to be identified, because they constrain which scenarios are plausible.

Formally, the structural scenario space is specified by a tuple $\langle N, A, J \rangle$ where N denotes a set of nodes, $A \subset N$ denotes a set of assumption nodes and J is a mapping that associates with each node $n \in N$ a set $J(n)$ of justifications. Each justification in $J(n)$ is itself a set of nodes corresponding to a combination of variables influencing n . This is illustrated by the sample hypergraph shown in Figure 4 with the associated specification of $\langle N, A, J \rangle$.

The structural scenario space is generated by means of two conventional inference techniques: (i) abduction of the plausible causes of known or hypothesised situations, and (ii) deduction of the plausible consequences of known or hypothesised situations. The abduction and deduction operations are applied to the emerging scenario space, by adapting the rules from the knowledge base to the specific circumstances described in the emerging scenario space:

- *Abduction*: given a piece of information that matches the consequent of a scenario fragment in the knowledge base, the abduction operation instantiates the information of the antecedents and assumptions of the rule, and adds them and the corresponding implication to the emerging scenario space. Consider, for example, the scenario fragment:

```

if {
  anesthetic-substance(A) }
assuming {
  is-addicted-to(P,A),
  has-recently-used(P,A) }
then {
  carries-traces-in-blood-of(P,A) }
distribution { ... }

```

And, assume that the scenario space contains a variable instance that indicates that johndoe has traces of percocet in his blood: i.e.

```
carries-traces-in-blood-of(johndoe, percocet)
```

This situation matches the consequent of the aforementioned scenario fragment with a substitution $\sigma = \{A/\text{percocet}, P/\text{johndoe}\}$. By application of abduction to the aforementioned scenario fragment, the antecedent `anesthetic-substance(percocet)`, and the assumptions `is-addicted-to(johndoe, percocet)` and `has-recently-used(johndoe, percocet)` are generated, the implication

```
anesthetic-substance(percocet) ∧
is-addicted-to(johndoe, percocet) ∧
has-recently-used(johndoe, percocet) → carries-traces-in-blood-of(johndoe, percocet)
```

is added to the emerging scenario space.

Note that the antecedent of certain scenario fragments may contain variables that are not contained in the consequent. For example, the scenario fragment

```
if {
  commits-homicide-by(P,V,M) }
then {
  died-from(V,M) }
distribution { ... }
```

states that if (a person) *P* successfully commits homicide of (person) *V* by means of (action) *M*, then it can be said that (person) *V* died from (action) *M*. Here, the antecedent contains variable *P*, which is not part of the consequent of the scenario fragment. Abduction of such rules requires that new instance names are created for those variables that the antecedent introduces. Following the Prolog convention, such instances will be given a name of the form *_n*, where *n* is an automatically generated integer: e.g. *_1*.

- *Deduction*: given a set of pieces of information that match the antecedents of a scenario fragment in the knowledge base, the deduction operation instantiates the information of the assumptions and consequent of the rule, and adds them and the corresponding implication to the emerging scenario space. Consider, for example, the following scenario fragment:

```
if {
  investigator(I),
  has-symptom(L,S),
  petechiae(S) }
assuming {
  correct-diagnosis(I,S) }
then {
  has-petechiae(L) }
distribution { ... }
```

And, assume that investigator `watson` makes an observation (of purple spots), herein identified as `_1`, on the eyes of `johndoe` and that that observation corresponds to `petechiae`: i.e.

```
investigator(watson),
has-symptom(eyes(johndoe),_1),
petechiae(_1) }
```

This situation matches the consequent of the aforementioned scenario fragment with a substitution $\sigma = \{I/watson, L/eyes(johndoe), S/_1\}$. By application of deduction to the above scenario fragment, the assumption `correct-diagnosis(watson,_1,is-addicted-to(johndoe,percocet))` and the piece of evidence `has-petechiae(eyes(johndoe))` are generated and the implication

```
investigator(watson) ∧ has-symptom(eyes(johndoe),_1) ∧
petechiae(_1) ∧ correct-diagnosis(watson,_1) → has-petechiae(eyes(johndoe))
```

is added to the emerging scenario space.

A formal specification of the procedure that generates structural scenario spaces is given by Algorithm 5.1. That is, `CONSTRUCTSCENARIOSPACESTRUCTURE(K, I, S, F)` synthesises a structural scenario space $\langle N, A, J \rangle$ for a given set of scenario fragments K , a set of established inconsistencies I , a set S of available evidence and a set F of known facts.

Algorithm 5.1: `CONSTRUCTSCENARIOSPACESTRUCTURE(K, I, S, F)`

```
 $\langle N, A, J \rangle \leftarrow \text{INITIALISESCENARIOSPACE}(S, F);$ 
ABDUCECAUSES( $\langle N, A, J \rangle, K$ );
DEDUCECONSEQUENCES( $\langle N, A, J \rangle, K$ );
ABDUCECAUSES( $\langle N, A, J \rangle, K$ );
DEDUCECONSEQUENCES( $\langle N, A, J \rangle, I$ );
SPURIOUSNODEREMOVAL( $\langle N, A, J \rangle$ );
return ( $\langle N, A, J \rangle$ );
```

Algorithm 5.1 consists of 6 steps:

1. *Initialisation:* A new symbolic scenario space is first created containing only the known facts and pieces of evidence and no inferences between them. More precisely, Algorithm 5.1 first initialises the hypergraph with the procedure `INITIALISESCENARIOSPACE(S, F)`, which is specified as Algorithm 5.2. This results in a hypergraph $S \cup F, F, J$, with the following justifications:

$$J(n) = \begin{cases} \{\} & \text{if } n \in S \\ \{\{\}\} & \text{if } n \in F \end{cases}$$

In other words, the pieces of evidence receive no justification yet and the facts are justified by the empty set because they are deemed true under all circumstances.

Algorithm 5.2: INITIALISESCENARIOSPACE(S, F)

```

 $N \leftarrow$  new set;
 $A \leftarrow$  new set;
 $J \leftarrow$  new table;
for each  $p \in S$ ,  $N \leftarrow N \cup \{p\}$ ;
for each  $p \in F$ 
  do  $\left\{ \begin{array}{l} N \leftarrow N \cup \{p\}; \\ A \leftarrow A \cup \{p\}; \end{array} \right.$ 
return ( $\langle N, A, J \rangle$ );

```

The procedure will be illustrated using the knowledge base of Appendix A . For simplicity, a situation with only a single piece of evidence will be assumed: the hanging body of man named johndoe. In this case, the initialisation phase produces a partial structural scenario space with only one node labelled evidence (hanging (body (johndoe))).

2. *Abduction of plausible causes from collected evidence:* The abduction operation explained above is applied with all the scenario fragments in the knowledge base whose consequent matches a situation variable instance in the emerging scenario space. In this way, causal chains of situations that constitute plausible explanations for the available evidence are constructed.

Algorithm 5.3: ABDUCECAUSES($\langle N, A, J \rangle, \mathbf{K}$)

```

for each substitution( $\sigma$ ),
   $\langle \text{if } P_{\text{antecedent}} \text{ assuming } P_{\text{assumptions}} \text{ then } \{p_c\} \text{distribution } D \rangle \in \mathbf{K},$ 
   $\sigma p_n \in N$ 
  do  $\left\{ \begin{array}{l} E \leftarrow \text{new set}; \\ \textbf{for each } p_i \in P_{\text{antecedent}} \cup P_{\text{assumptions}} \\ \textbf{do} \left\{ \begin{array}{l} \textbf{if } \sigma p_i \notin N \\ \textbf{do} \left\{ \begin{array}{l} p \leftarrow \text{instantiate}(\sigma p_i); \\ N \leftarrow N \cup \{p\}; \\ \textbf{if } p_i \in P_{\text{assumptions}} \\ \textbf{then } A \leftarrow A \cup \{p\}; \\ E \leftarrow E \cup \{p\}; \\ \textbf{else } E \leftarrow E \cup \{\sigma p_i\}; \end{array} \right. \end{array} \right. \\ J(\sigma p_c) \leftarrow J(\sigma p_c) \cup \{E\}; \end{array} \right.$ 

```




Figure 5: Hypergraph generated after applying Algorithm 5.3 in Step 2. Note: the shaded boxes represent new nodes added in this step of the procedure.

Formally, Algorithm 5.3 generates these explanations for the available 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 .

Figure 5 depicts a sample partial scenario space that would emerge after this step if the knowledge base of Appendix A and the single piece of evidence introduced in Step 1 are employed.

3. *Deduction of plausible consequences from plausible causes:* The deduction operation explained above is applied with all the scenario fragments in the knowledge base whose antecedents match a set of situation variable instances in the emerging scenario space. In so doing, causal chains containing the consequences of those situations that are plausible causes of the available evidence are generated. The objective of this step in the overall procedure is to infer the evidence that has not yet been identified, but may be produced by plausible scenarios that explain the available evidence. As discussed below, it is important to identify these pieces of evidence because this information helps the system in supporting the investigator to determine which of the scenarios that explain the available evidence may be more plausible.

Algorithm 5.4: DEDUCECONSEQUENCES($\langle N, A, J \rangle \mathbf{K}$)

```

for each substitution( $\sigma$ ),
   $\langle \text{if } P_{\text{antecedent}} \text{ assuming } P_{\text{assumptions}} \text{ then } \{p_c\} \text{distribution } D \rangle \in \mathbf{K},$ 
   $\{\sigma p_i \mid p_i \in P_{\text{antecedent}}\} \subset N$ 
  do {
    if  $\sigma p_c \notin N$ 
    then {
       $p' \leftarrow \text{instantiate}(\sigma p_c);$ 
       $N \leftarrow N \cup \{p'\};$ 
       $E \leftarrow \text{new set};$ 
      for each  $p_i \in P_{\text{antecedent}}$ 
      do  $E \leftarrow E \cup \{\sigma p_i\};$ 
      for each  $p_i \in P_{\text{assumptions}}$ 
      do {
        if  $\sigma p_i \notin N$ 
        then {
           $p \leftarrow \text{instantiate}(\sigma p_i);$ 
           $N \leftarrow N \cup \{p\};$ 
           $A \leftarrow A \cup \{p\};$ 
           $E \leftarrow E \cup \{p\};$ 
        }
        else  $E \leftarrow E \cup \{\sigma p_i\};$ 
      }
       $J(p') \leftarrow J(p') \cup \{E\};$ 
    }
  }

```

Algorithm 5.4 generates all possible consequences of the explanations created previously. 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.

Continuing with the ongoing example, Figure 6 depicts a sample partial scenario space that would emerge after this step.

4. *Abduction of plausible causes from uncollected evidence:* When assessing the value of searching for additional evidence through further investigative actions, it is important that the entire range of possible causes of such evidence is considered, not just the causes that reside within the crime scenarios. For example, when searching for gun powder residue on the hands of a suspect, the possible role of this suspect in the crime under investigation is not the only relevant plausible explanation for a potential discovery of gun powder residue. The background of the suspect may constitute an alternative explanation when, for instance, the suspect is an amateur hunter. Therefore, Step 2 of the procedure is repeated for the pieces of uncollected, but potentially available evidence, which have been deduced in Step 3.

The partial scenario space that results from this step of the structural scenario space synthesis algorithm is shown in

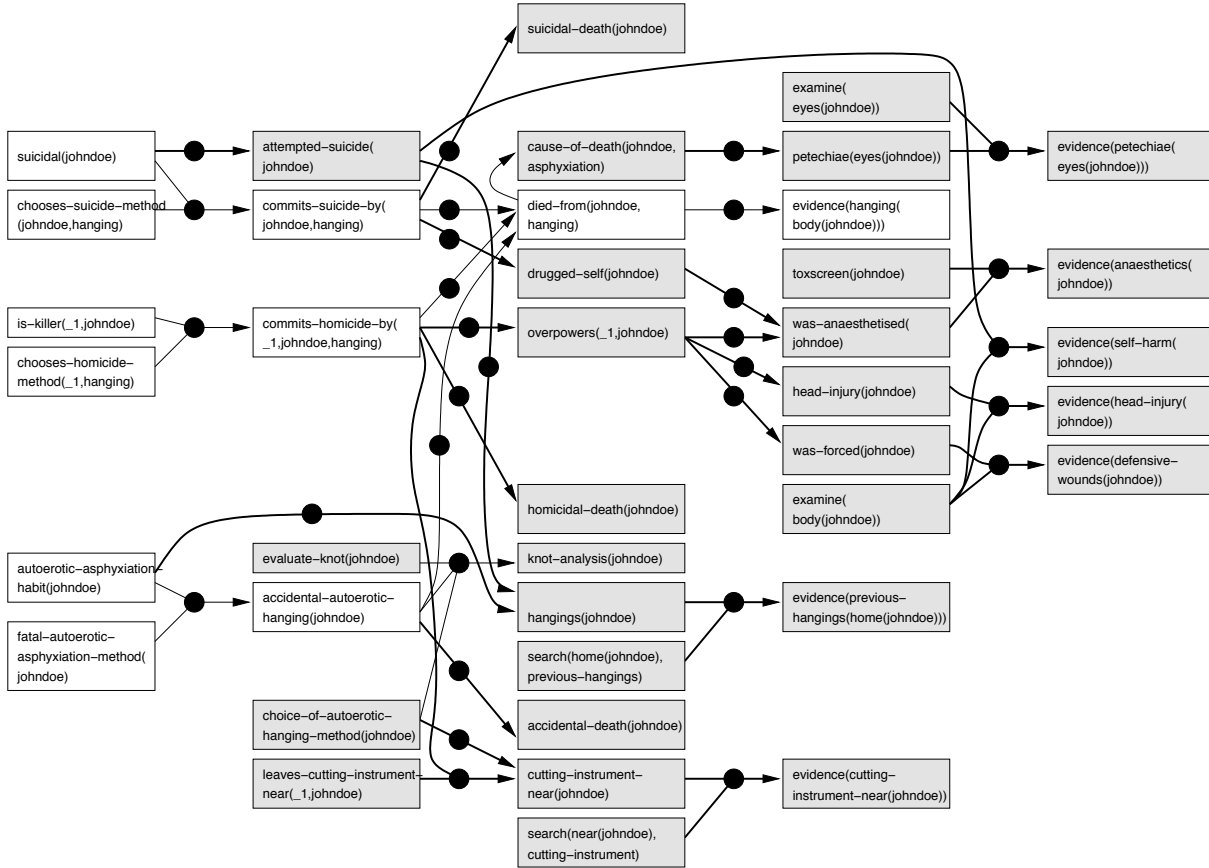


Figure 6: Hypergraph generated after applying Algorithm 5.4 in Step 3. Note: the shaded boxes represent new nodes added in this step of the procedure.




Figure 7: Hypergraph generated after applying Algorithm 5.3 in Step 4. Note: the shaded boxes represent new nodes added in this step of the procedure.

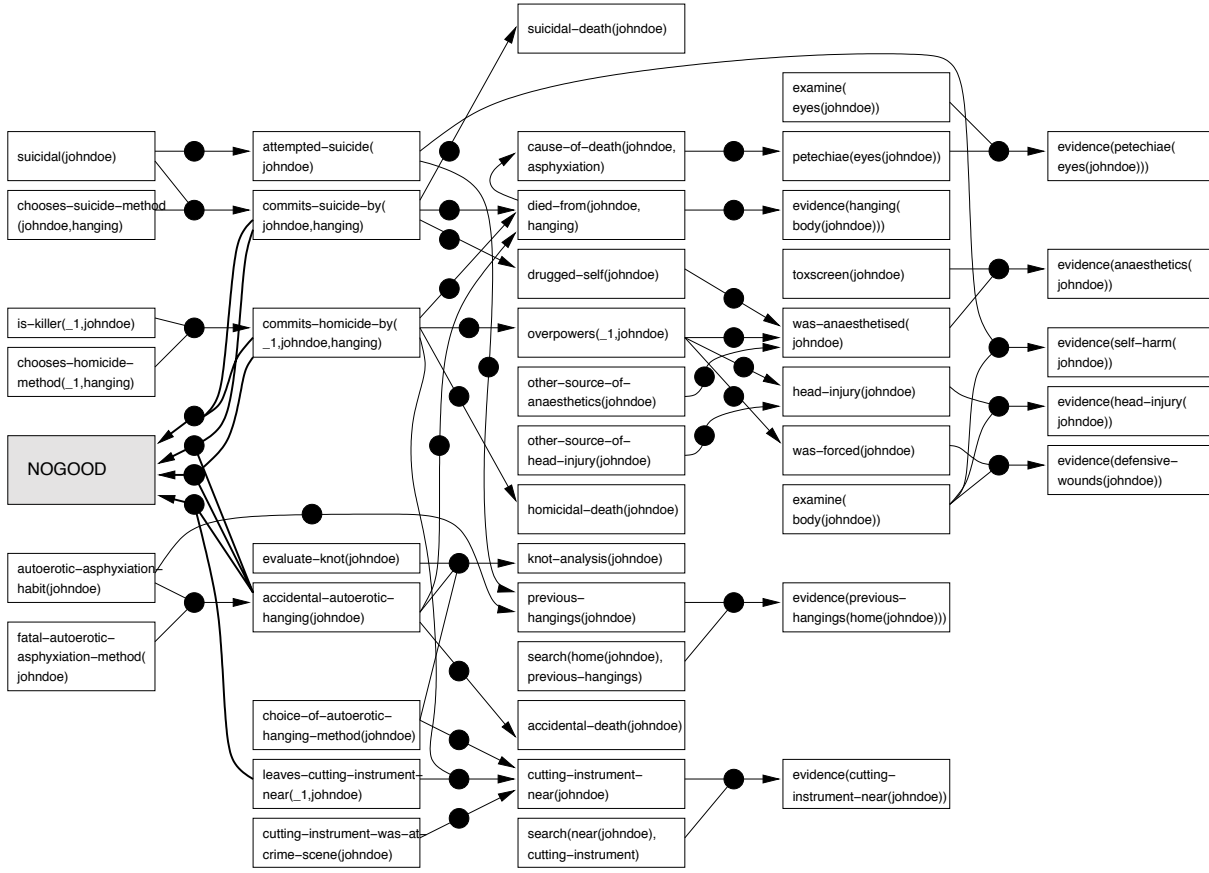


Figure 8: Hypergraph generated after deducing inconsistencies in Step 5

Figure 7.

5. *Deduction of inconsistencies*: Once a symbolic scenario space has been constructed that contains the available evidence (Step 1), the possible causes of the available evidence (Step 2), the possible but uncollected evidence that may be generated by these possible causes (Step 3) and the possible causes of the uncollected evidence (Step 4), any given constraints under which situations that may be part of the same scenario are added to the emerging symbolic scenario space. This involves applying a deduction operation for each inconsistency whose situation variables match a set of situation variables in the emerging symbolic scenario space.

The partial scenario space that results from this step of the structural scenario space synthesis algorithm is shown in Figure 8.

6. *Spurious explanation removal*: The structural scenario space generated by the above procedure enables the decision support system to trace back the conjectures and facts that could originally have caused the available evidence. Because such causal chains may have to be established by multiple abduction operations applying different scenario fragments, it is possible that the resulting structural scenario space contains situations that are not a consequence of any situation and that are not assumptions or facts. As explained in section 4, 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 probabilistic scenario space as they represent the only information which is associated with a prior distribution. Therefore, such nodes constitute spurious explanations, and they are to be removed from the emerging scenario space.

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.

Algorithm 5.5: SPURIOUSNODEREMOVAL($\langle N, A, J \rangle$)

```

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

```

As the hypergraph of Figure 8 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) : T` is a fact. Furthermore, `suspect(johndoe)` and `suspect(_1)` correspond to assumption nodes, where `suspect(johndoe) : T` 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.

5.2 Generating the probabilistic scenario space

As explained above, the probabilistic scenario space is a Bayesian network (BN) that represents the extent to which the situations in the structural scenario space affect one another. Such a scenario space is created by following the procedure

below:

- For each situation variable instance in the structural scenario space, create a corresponding node;
- For each implication $\dots \wedge a \wedge \dots \rightarrow c$, where a and c refer to situation variable instances, create an arc $a \rightarrow c$;
- For each assumption or fact in the structural scenario space, which (as defined above) is always represented by a parent node and not by a child node, obtain a prior probability distribution, which must either be predefined in the knowledge base or be specified by the user; and
- For each node that does not correspond to an assumption or fact, calculate a conditional probability distribution table based on the probability distributions associated with each scenario fragment.

The first three steps of this procedure are straightforward and are formally specified as follows:

Algorithm 5.6: $\text{CREATEDAG}(\langle N, A, J \rangle)$

$G \leftarrow \text{new DAG};$

for each $n \in N$

do $\begin{cases} \text{add}(G, n); \\ \text{for each } n' \in (\bigcup_{E \in J(n)} E) \text{ } \mathbf{return} (G); \\ \quad \mathbf{do} \text{ add}(G, \text{arc}(n', n)); \end{cases}$

The generation of the conditional probability tables is somewhat more complex. Therefore, the remainder of this subsection is dedicated to this topic.

Let $\{p_1, \dots, p_s\}$ be the set of parents of p_n in the probabilistic scenario space, $p_c : c$ be an assignment of p_c , and A be a set of assignments $\{p_1 : v_1, \dots, p_s : v_s\}$, where each $v_i \in D_{p_i}$. The conditional probability of the situation $p_c : c$ given A is determined by those scenario fragments that determine how the situations $p_1 : v_1, \dots, p_s : v_s$ affect p_c . Let SF_1, \dots, SF_k be such scenario fragments and $p_c : c_1, \dots, p_c : c_k$ be their respective outcomes. p_c is assigned c whenever the combined effect of the scenario fragments $c_1 \oplus \dots \oplus c_k = c$, with regard to a certain predefined interpretation of the combination operator \oplus . From this, the conditional probability $P(p_c : c \mid A)$ is given by:

$$P(p_n : c \mid A) = P \left[\bigvee_{c_1 \oplus \dots \oplus c_k = c} \left(\bigwedge_{i=1, \dots, k} (A \xrightarrow{SF_i} p_n : c_i) \right) \right] \quad (2)$$

Thus, computing $P(p_n : c \mid A)$ involves calculating the likelihood of a combination of events described by a disjunctive normal form (DNF) expression (i.e. an expression of the form $(x_{11} \text{ and } \dots \text{ and } x_{1m_1}) \text{ or } \dots \text{ or } (x_{n1} \text{ and } \dots \text{ and } x_{nm_n})$).

Because the occurrence of the different combinations of outcomes c_1, \dots, c_k of scenario fragments SF_1, \dots, SF_k involves mutually exclusive events, the calculation can be resolved by adding the probabilities of the conjuncts in (2):

$$P(p_n : c \mid A) = \sum_{c_1 \oplus \dots \oplus c_k = c} P \left(\bigwedge_{i=1, \dots, k} (A \xrightarrow{SF_i} p_n : c_i) \right) \quad (3)$$

From presumption 1 (as noted in 4.3), the outcomes of different scenario fragments (of the same consequent), with regard to a given set of assignments of the antecedent and assumption variables, correspond to independent events. Therefore, the probability of the conjunctions in (3) is equal to the product of the probabilities of their conjuncts, thereby

$$P(p_n : c \mid A) = \sum_{c_1 \oplus \dots \oplus c_k = c} \left(\prod_{i=1, \dots, k} P(A \xrightarrow{SF_i} p_n : c_i) \right) \quad (4)$$

Consider the following two scenario fragments, which form part of the probabilistic scenario space that corresponds to the structural space of Figure 8:

```

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

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

```

where `autoerotic-asphyxiation-habit(V)` and `attempted-suicide(V)` correspond to boolean variables, and `previous-hangings(V)` to a variable taking values from the (totally ordered) domain `{never, veryfew, several}`, with `never < veryfew < several`. The combination operator associated with the latter variable is presumed to be `max`.

For notational convenience, let the above two scenario fragments be denoted by SF_1 and SF_2 respectively, let `autoerotic-asphyxiation-habit(V)`, `attempted-suicide(V)` and `previous-hangings(V)` be represented as p_1 , p_2 and p_3 respectively and let \top and \perp symbolise true and false. This permits the probability distributions of the scenario fragments to be expressed as follows:

$$\begin{aligned}
P(p_1 : \top \xrightarrow{SF_1} p_3 : \text{never}) &= 0.10 \\
P(p_1 : \top \xrightarrow{SF_1} p_3 : \text{veryfew}) &= 0.40 \\
P(p_1 : \top \xrightarrow{SF_1} p_3 : \text{several}) &= 0.50 \\
P(p_2 : \top \xrightarrow{SF_2} p_3 : \text{never}) &= 0.70 \\
P(p_2 : \top \xrightarrow{SF_2} p_3 : \text{veryfew}) &= 0.29 \\
P(p_2 : \top \xrightarrow{SF_2} p_3 : \text{several}) &= 0.01
\end{aligned}$$

Then, the probabilities of assignments to previous-hangings (V), given that autoerotic-asphyxiation-habit (V) and attempted-suicide (V) are assigned true (or the probabilities of assignments to p_3 given $p_1 : \top$ and $p_2 : \top$), can be computed as follows.

According to (4), all combinations of the outcomes c_1 and c_2 of scenario fragments SF_1 and SF_2 , with $c_1, c_2 \in \{\text{never}, \text{veryfew}, \text{several}\}$, such that $\max(c_1, c_2) = \text{veryfew}$, must be considered. There are three such combinations: $\{c_1 : \text{veryfew}, c_2 : \text{veryfew}\}$, $\{c_1 : \text{never}, c_2 : \text{veryfew}\}$ and $\{c_1 : \text{veryfew}, c_2 : \text{never}\}$. Hence, $P(p_3 : \text{veryfew} | p_1 : \top, p_2 : \top)$ can be computed by

$$\begin{aligned}
&P(p_3 : \text{veryfew} | p_1 : \top, p_2 : \top) \\
&= P(p_1 : \top \xrightarrow{SF_1} p_3 : \text{veryfew}) \times P(p_2 : \top \xrightarrow{SF_2} p_3 : \text{veryfew}) + \\
&\quad P(p_1 : \top \xrightarrow{SF_1} p_3 : \text{never}) \times P(p_2 : \top \xrightarrow{SF_2} p_3 : \text{veryfew}) + \\
&\quad P(p_1 : \top \xrightarrow{SF_1} p_3 : \text{veryfew}) \times P(p_2 : \top \xrightarrow{SF_2} p_3 : \text{never}) \\
&= 0.4 \times 0.29 + 0.1 \times 0.29 + 0.4 \times 0.7 = 0.425
\end{aligned}$$

Similarly, it can be shown that

$$\begin{aligned}
P(p_3 : \text{never} | p_1 : \top, p_2 : \top) &= 0.070 \\
P(p_3 : \text{several} | p_1 : \top, p_2 : \top) &= 0.505
\end{aligned}$$

The complete scenario space for the ongoing example is a Bayesian network containing 43 variables, 28 conditional probability tables and 15 prior probability distributions. Clearly, any representation of the full specification of this network, and the computations required for belief propagation by means of this network, would significantly add to the size of this paper while contributing little new knowledge. Thus, only a part of this BN is shown in Figure 9 for illustration, which relates to the corresponding part of the scenario space of Figure 5.

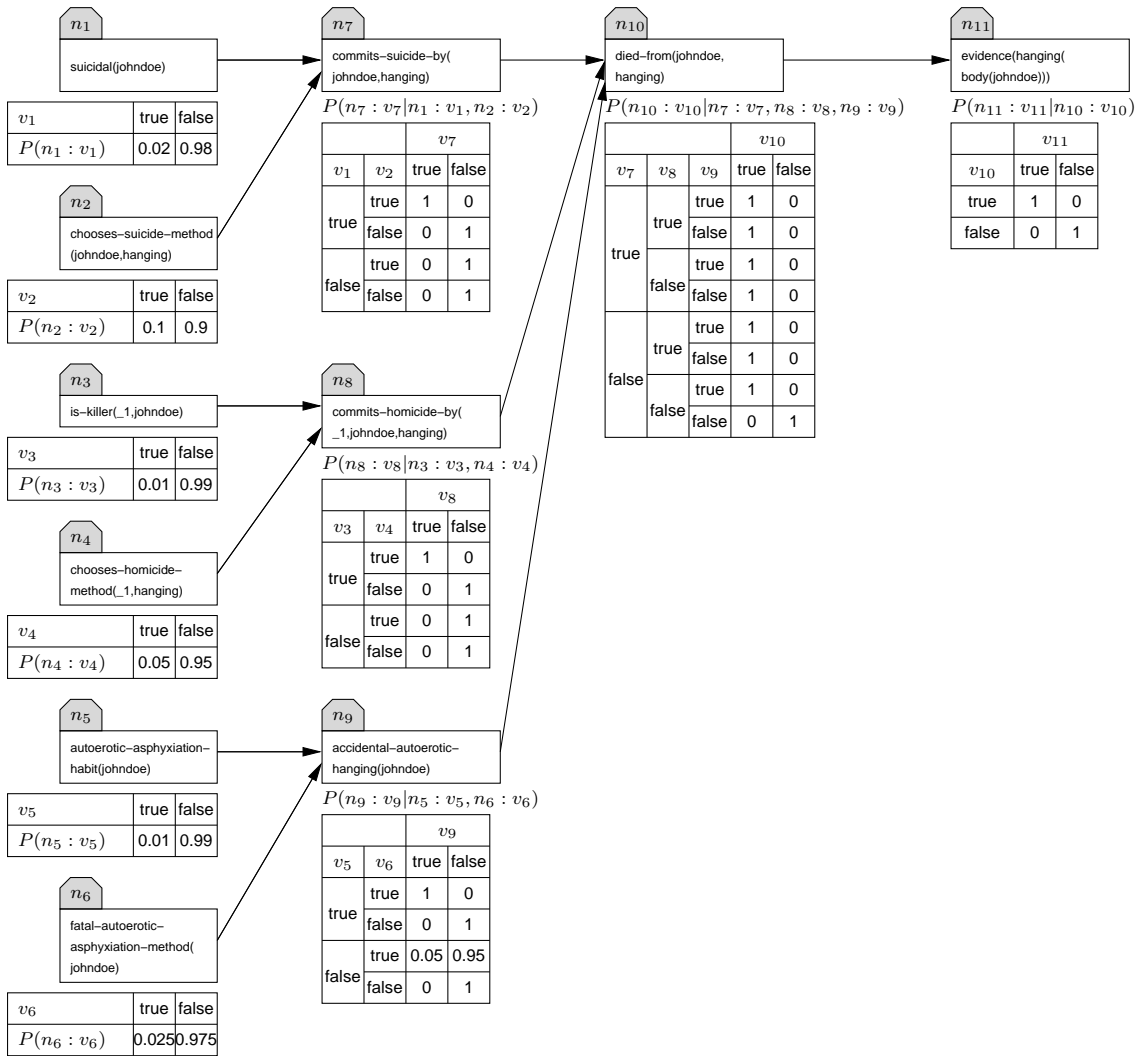


Figure 9: Sample partial Bayesian network

6 Scenario Space Analysis

Once constructed, the probabilistic scenario space can be analysed in conjunction with the structural one to compute effective evidence collection strategies. The concepts of evidence, hypotheses, assumptions and facts are still employed in the probabilistic 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 can be extended).

The work will be illustrated by means of probabilities derived from a BN which has been generated by means of the techniques of section 5 using the knowledge base of Appendix A (and which is a Bayesian representation of the symbolic scenario space given in Figure 8).

6.1 Hypothesis sets and query types

The approach shown here aims to interpret the investigatory worth of different scenarios in relation to a set H of hypotheses. This set must be exhaustive and the hypotheses within it mutually exclusive. The set H is *exhaustive* if one of the hypotheses in the set is guaranteed to be true. This ensures that the approach will evaluate the scenario space entirely. The hypotheses in a set are *mutually exclusive* if no pair of hypotheses taken from the set can be true simultaneously. This simplifies the interpretation of entropy values because it ensures that a situation of total certainty corresponds to an entropy value of 0. These conditions may seem to be rather strong in general, but they are not strong for the present work. This is because the exhaustivity is only subject to the demand for the generated space to be fully examined in the evaluation, while different scenarios naturally satisfy the requirement that only one underlying scenario be true at any one time. An example of a hypothesis set that meets these criteria is:

$$H = \{\text{death of johndoe was homicide} : \top, \\ \text{death of johndoe was suicide} : \top, \\ \text{death of johndoe was accidental} : \top, \\ \text{death of johndoe was natural} : \top\}$$

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 cause of death of a 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:

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

It is the responsibility of the knowledge engineer to ensure that the hypothesis sets generated in this way meet the exhaustiveness and mutual exclusivity criteria. These criteria can be satisfied for any given set $P = \{p_1, \dots, p_n\}$ of predicates identifying hypothesis 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 $p_1 : \perp, \dots, p_n : \perp$, and $p_{n+1} : \perp$ with likelihood 1 otherwise:

```

if {p1, ..., pn}
then {pn+1}
distribution pn+1 {⊥, ..., ⊥ → ⊤ : 1, ⊥ : 0}

```

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

```

inconsistent {pi : ⊤, pj : ⊤}

```


6.2 Entropy

As indicated previously, the work here employs an information theory based approach, which is widely used in areas such as machine learning [39] and model based diagnosis [23]. Information theory utilises the concept of *entropy*, which is a measurement of doubt over a range of choices. Applied to the present problem, the entropy over an exhaustive set of mutually exclusive hypotheses $H = \{h_1, \dots, h_m\}$ is given by:

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

where the log base is 2, and the values $P(h)$ can be computed by conventional BN inference techniques. Intuitively, entropy can be interpreted as a quantification of a 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 $P(h_1) = \dots = P(h_m) = \frac{1}{m}$ and $\epsilon(H)$ reaches 0, its lowest value (which corresponds to a totally certain situation) when all $P(h_i) = 0, i = 1, \dots, j-1, j+1, \dots, m$ and $P(h_j) = 1, j \in \{1, \dots, m\}$.

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

$$\epsilon(H \mid E) = - \sum_{h \in H} P(h \mid E) \log P(h \mid E) \quad (5)$$

where the values $P(h \mid E)$ can, again, be computed through BN inference. This expression may be considered as a measure of the lack of information H given E . For the example problem from the sample scenario space, the following probabilities can be computed, with E_1 containing evidence `(hanging (body (johndoe))) : T` and `nogood : ⊥`:

$$P(\text{homicidal-death}(\text{johndoe}) \mid E_1) = 0.22$$

$$P(\text{suicidal-death}(\text{johndoe}) \mid E_1) = 0.33$$

$$P(\text{accidental-death}(\text{johndoe}) \mid E_1) = 0.45$$

Thus, as an instance,

$$P(H_1 \mid E_1) = -(0.22 \log 0.22 + 0.33 \log 0.33 + 0.45 \log 0.45) = 1.53$$

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

$$\min_{a \in A} Exp(\epsilon(H \mid E), a) \quad (6)$$

Note that the entropy values calculated by (5) are affected by the prior distributions assigned to the assumptions, as described in section 4. Within the context of evidence evaluation (which often applies the likelihood ratio based approach [1]), 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. 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, the strategy dictated by (6) offers a useful means to decide on what evidence to collect next. Nevertheless, the minimal entropy decision rule does not yield information that should be used for evidence evaluation in court.

6.3 Minimal entropy-based evidence collection

Let a denote an investigative action and E_a be a set of the variable assignments corresponding to the 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 with regard to the different possible outcomes $e \in E_a$, weighted by the likelihood of obtaining each outcome e (given the available evidence):

$$Exp(\epsilon(H \mid E), a) = \sum_{e \in E_a} P(e \mid a : \top, E) \epsilon(H \mid E \cup \{a : \top, e\}) \quad (7)$$

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

$$P(\text{toxscreen}(\text{johndoe}) : \top \mid E_2) = 0.17$$

$$P(\text{toxscreen}(\text{johndoe}) : \perp \mid E_2) = 0.83$$

$$P(\text{homicidal-death}(\text{johndoe}) \mid E_2 \cup \{\text{toxscreen}(\text{johndoe}) : \top\}) = 0.40$$

$$P(\text{suicidal-death}(\text{johndoe}) \mid E_2 \cup \{\text{toxscreen}(\text{johndoe}) : \top\}) = 0.49$$

$$P(\text{accidental-death}(\text{johndoe}) \mid E_2 \cup \{\text{toxscreen}(\text{johndoe}) : \top\}) = 0.11$$

$$P(\text{homicidal-death}(\text{johndoe}) \mid E_2 \cup \{\text{toxscreen}(\text{johndoe}) : \perp\}) = 0.19$$

$$P(\text{suicidal-death}(\text{johndoe}) \mid E_2 \cup \{\text{toxscreen}(\text{johndoe}) : \perp\}) = 0.44$$

$$P(\text{accidental-death}(\text{johndoe}) \mid E_2 \cup \{\text{toxscreen}(\text{johndoe}) : \perp\}) = 0.38$$

Intuitively, these probabilities can be explained as follows. In a homicide, anaesthetics may have been used by the murderer to gain control over johndoe, and in a suicide case, johndoe may have used anaesthetics during 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:

$$\begin{aligned} \text{Exp}(\epsilon(H \mid E_1), a) &= 0.17 \times -(0.40 \log 0.40 + 0.49 \log 0.49 + 0.11 \log 0.11) + \\ &\quad 0.83 \times -(0.19 \log 0.19 + 0.44 \log 0.44 + 0.38 \log 0.38) \\ &= 0.17 \times 1.38 + 0.83 \times 1.51 = 1.49 \end{aligned}$$

The investigative action that is expected to provide the most information is the one that minimises the corresponding EPE. For example, Table 2 shows a number of possible investigative actions that can be undertaken (in column 1) and the corresponding EPEs in the sample probabilistic 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 therefore established to be 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, which have both been the most likely hypotheses if anaesthetics are discovered in the body.

6.4 Extensions

While the approach presented above is itself a useful alternative to the likelihood ratio approach, several further improvements of this proposed technique are proposed. As these extensions are not difficult to conceive, detailed examples are

omitted below.

6.4.1 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 such one-step lookahead optimisation approach. The likelihood of obtaining low quality locally optimal evidence collection strategies can be reduced by considering the EPEs after performing a sequence of actions a_1, \dots, a_v (of course, with incurred overheads over computation):

$$\begin{aligned}
 &Exp(\epsilon(H \mid E), a_1, \dots, a_v) \\
 &= \sum_{e_1 \in E_{a_1}} \dots \sum_{e_v \in E_{a_v}} P(e_1, \dots, e_v \mid a_1 : \top, \dots, a_v : \top, E) \\
 &\quad \epsilon(H \mid e_1, a_1 : \top, \dots, e_v, a_v : \top, E)
 \end{aligned} \tag{8}$$

In order to determine $Exp(\epsilon(H \mid E), a_1, \dots, a_v)$, equation (8) can be simplified as follows:

$$\begin{aligned}
 &Exp(\epsilon(H \mid E), a_1, \dots, a_v) \\
 &= \sum_{e_1 \in E_{a_1}} \dots \sum_{e_v \in E_{a_v}} \frac{P(e_1, \dots, e_v, a_1 : \top, \dots, a_v : \top, E)}{a_1 : \top, \dots, a_v : \top, E} \\
 &\quad \epsilon(H \mid E \cup \{e_1, a_1 : \top, \dots, e_v, a_v : \top\}) \\
 &= \sum_{e_1 \in E_{a_1}} \dots \sum_{e_v \in E_{a_v}} \left(\prod_{i=1}^v \frac{P(e_i, a_i : \top, \dots, a_v : \top, E)}{a_1 : \top, \dots, a_v : \top, E} \right) \\
 &\quad \epsilon(H \mid E \cup \{e_1, a_1 : \top, \dots, e_v, a_v : \top\}) \\
 &= \sum_{e_1 \in E_{a_1}} \dots \sum_{e_v \in E_{a_v}} \left(\prod_{i=1}^v P(e_i \mid a_1 : \top, \dots, a_v : \top, E) \right) \\
 &\quad \epsilon(H \mid E \cup \{e_1, a_1 : \top, \dots, e_v, a_v : \top\})
 \end{aligned} \tag{9}$$

6.4.2 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 outcomes 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 $E_{a,1}, \dots, E_{a,w}$:

$$\begin{aligned}
& Exp(\epsilon(H \mid E), a) \\
&= \sum_{e_1 \in E_{a,1}} \dots \sum_{e_w \in E_{a,w}} P(e_1, \dots, e_w \mid a : \top, E) \\
&\quad \epsilon(H \mid e_1, \dots, e_w, a : \top, E) \\
&= \sum_{e_1 \in E_{a,1}} \dots \sum_{e_w \in E_{a,w}} \left(\prod_{i=1}^w P(e_i \mid a : \top, E) \right) \\
&\quad \epsilon(H \mid E \cup \{e_1, \dots, e_w, a : \top\})
\end{aligned} \tag{10}$$

6.4.3 Multiple hypothesis sets

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

$$\begin{aligned}
& \epsilon(H_1, \dots, H_t \mid E) \\
&= - \sum_{h_1 \in H_1} \dots \sum_{h_t \in H_t} P(h_1, \dots, h_t \mid E) \log P(h_1, \dots, h_t \mid E) \\
&= - \sum_{h_1 \in H_1} \dots \sum_{h_t \in H_t} \left(\prod_{i=1}^t P(h_i \mid E) \right) \log \left(\prod_{i=1}^t P(h_i \mid E) \right)
\end{aligned} \tag{11}$$

6.5 User interface

While a detailed discussion of the user interface developed for the present 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 realising the objectives of the system. Investigators may have a number of considerations that are beyond the scope of the current implementation. These may include perishability of evidence, legal restrictions, limitations on resources and overall workload. Therefore, the implemented system has been devised to list alternative evidence collection strategies in increasing order of EPEs.

Investigative action	EPE	NEER	REER
Knot analysis	0.98	26%	29%
Examine body	1.12	17%	19%
Search for cutting instrument	1.19	13%	14%
Search for signs of previous hangings	1.36	1.3%	1.5%
Check eyes for petechiae	1.38	0%	0%

Table 2: Evaluation of investigative actions

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 due to performing an investigative action a (i.e. $\epsilon(H \mid E) - \text{Exp}(\epsilon(H \mid E), a)$) as a proportion of the maximal entropy under total lack of information. As such, it provides a means of assessing case progress:

$$NEER(H \mid E, a) = \frac{\epsilon(H \mid E) - \text{Exp}(\epsilon(H \mid E), a)}{\epsilon(H)} \quad (12)$$

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

$$REER(H \mid E, a) = \frac{\epsilon(H \mid E) - \text{Exp}(\epsilon(H \mid E), a)}{\epsilon(H \mid E)} \quad (13)$$

These calculations are illustrated in Table 2 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.

7 Discussions

7.1 Related work

Evidential reasoning is an important, current topic in Artificial Intelligence and Law. A broad range of systems have been developed for this purpose. The system presented here is distinct from the existing literature in that it addresses a different but nonetheless significant problem in evidential reasoning: hypothesis synthesis and their analysis with a view to devising evidence collection strategies.

Argumentation systems [8, 38, 47, 55, 58] and probabilistic expert systems [2, 9, 40] are primarily intended to support evidential reasoning in courts. Most argumentation systems represent plausible causes as consequences of obtained evidence, in order to contrast alternate lines of reasoning [49]. One notable exception is AVER, which combines the aforementioned explanatory (observation-to-cause) inferences with causal (cause-to-effect) inferences [3]. AVER, however, aims to model investigators' narratives; it is not designed to support the calculation of effective evidence collection strategies. Probabilistic expert systems aim to calculate the extent to which a single piece of evidence supports one scenario rather than competing scenarios, with a view to justify the testimony of forensic experts in court [9]. Thus, these types of system are designed to be employed after the evidence collection effort that the DSS described here supports.

Data mining systems for intelligence databases [6, 30, 42, 60] and other types of data analysis tool [16, 24, 25, 48, 54] seek to identify possible crimes. While such tools determine whether there is evidence of elements of a certain crime scenario, they are not intended to reconstruct plausible alternatives of crime scenarios. In other words, these types of system support analysis of plausible crimes rather than plausible scenarios of a single crime. As such, they are designed to be employed before an individual case is examined in detail.

7.2 Scaling up applications

The DSS described herein incorporates the methodologies for scientific inquiry in crime investigation, following the best practice proposed by forensic scientists [9, 29]. A prototypical system has been developed, and it was evaluated in collaboration with forensic scientists at the Forensic Laboratory of a UK Police Force, with regard to a number of solicited knowledge bases. A small-scale example of such a knowledge base is shown in the appendix. This has demonstrated the DSS's ability to construct and simultaneously analyse a broad variety of plausible crime scenarios from compact knowledge bases. As such, the system is potentially able to support the construction and refinement of plausible scenarios during the early investigation of serious crimes. This is a novel development and distinguishes the DSS described herein from the other DSSs built in the domain of crime investigation.

The present work contains two key features that facilitate scaled up applications. Firstly, the compositional modelling approach allows causal knowledge concerning specific aspects of hypothetical situations and events, or specific types of evidence to be defined from first principles. First principles can be considered in isolation from the extraneous context, such as unrelated or partially related conjecture and evidence or any investigative procedure, in which it is used. Such knowledge is less complex and more reusable compared to experienced whole cases. More specifically, individual serious crime cases reoccur relatively infrequently compared to the constituent hypothetical events and evidence. As such, the use of a knowledge base relying on first principles makes the approach more robust. Similar observations have been made in the domain of physical systems diagnosis [7].

Secondly, the maximum entropy reduction approach, when applied to the model space, is capable of considering a large number of plausible scenarios and relating them to available evidence and investigative activities. It has the ability to

identify the potentially most productive evidence collection strategies. This is beneficial as simultaneous analysis of multiple scenarios, in order to guide an ongoing investigation, is an area in which both human investigators and conventional decision support systems lack. Despite these advantages, difficulties in deploying such a system remain. These issues and ways in which these difficulties may be resolved are discussed in the remainder of this section.

7.3 First principles and the *ceteris paribus* assumption

As indicated previously, the approach presented herein builds on model based reasoning techniques that have been applied successfully in the physical domain. Of course, there are crucial differences between the natural science, which studies physical systems, and forensic science, which studies crime scenarios. The most important of these is that the behaviour of physical systems can be described by first principles that can be applied *ceteris paribus* (i.e. under the assumption of all other things being equal). While there are certain first principles in the forensic science domain, the *ceteris paribus* clause generally does not hold. A key reason for this is that crime scenes are not pristine, uncontaminated places [41]. Consequently, there are many factors that may affect the way in which combinations of causes influence (observable) consequences. For example, the amount of blood splatter transferred from the victim to the perpetrator due to an assault with a blunt instrument is proportional to the extent and duration of the assault. There are many factors that may make subsequent discovery of blood splatter on the perpetrator matching the victim more or less likely (e.g. the propensity of the victim to bleed as a consequence of injury, the existence of blood from sources other than the victim on the perpetrator's clothes, etc.). Therefore, the presumption of independence of any two probability distributions taken from two scenario fragments involving the same consequent variable is difficult to uphold.

7.4 Subjective probability

Obviously, first principles in the forensic science domain are often a matter of approximation. Many factors need to be considered in determining the effect of an influence (e.g. the number of fibres transferred between two people involved in a physical altercation). It is extremely difficult, if not impossible, to model such factors precisely in practice. The use of probabilities may allow for uncertainty which is associated with the effects resulting from unknown or unconsidered factors to be expressed. However, probability distributions over scenario fragments usually stem from subjective assessments by experts rather than from experimental estimation. Therefore, the use of numeric probabilities conveys an inappropriate degree of precision [22].

This issue can be examined in more detail by considering a thought experiment as proposed in [56]. In such an experiment, a game is played between two experts using the following procedure. First, one expert offers a price q for a promise to pay 1 currency unit if a proposition of interest, denoted by $p : v$ here, turns out to be correct. Next, the other expert decides whether the first one may buy or sell such a promise. Finally, the actual proposition is revealed. Thus, q is a price at which the first expert is neutral between buying or selling a promise to pay 1 currency unit if $p : v$ is correct. The outcome

of the game is that, according to the first expert, $P(p : v) = q$. As such, this thought experiment provides a way to attach semantics to the subjective probabilities that are suggested by experts. In practice, however, experts are reluctant to commit to such numerical values. This implies that there is no a single obvious price at which experts are neutral between buying and selling a promise of the aforementioned kind.

There exists a range of extensions on classical Pascalian probability theory that facilitate capturing subjective probabilities by means of this thought experiment. These approaches are based on the use of intervals of probabilities, rather than the conventional point probabilities. An interval offers a richer representation to capture an expert's belief. In terms of the thought experiment of [56], the lower boundary of a probability interval, associated with a particular proposition and suggested by an expert, corresponds to the highest price at which the expert is willing to *buy* a promise to pay 1 currency unit if the proposition is revealed to be true. Similarly, the upper boundary of the probability interval corresponds to the lowest price at which the expert is willing to *sell* a promise to pay 1 currency unit if the proposition is revealed to be true.

Approaches that use probability intervals include qualitative probability networks and qualitative certainty networks [4, 15, 50], and also the order-of-magnitude based extensions thereof [43, 44]. Recent work has shown that QPNs with order-of-magnitude calculus can be adapted to express second-order influences (i.e. influences that affect the effect of other influences) [31]. Such second-order influences include, for example, inhibitors that reduce the effect of another influence, and amplifiers that increase the effect of another influence. Thus, they provide a mechanism to represent, in the form of scenario fragments, considerations that may sway an influence's outcome from its *ceteris paribus* effect. Furthermore, the influences in QPNs are easily compoundable into more complex models without the need for unrealistic assumptions, and they are intrinsically approximate in nature. As such, this line of research has great potential to address the aforementioned concerns.

More sophisticated variations on the interval based approach employ sets of intervals. These include second-order probability theory [11, 20, 57] and linguistic probability theory [22]. Such approaches relate probability intervals to another metric concerning each individual interval. This may be an expression of confidence in the interval, as in second-order probability theory, or a degree of belonging described by means of fuzzy membership values, as in linguistic probability theory.

7.5 Domain knowledge acquisition

In this work, the DSS' knowledge base is presumed to contain generic causal relations. These relations formalise how crime scenario events and situations have observable consequences that may be collected and analysed to form evidence. Analogously, typical forensic science textbooks describe types of evidence, the methodology for their identification, collection and preservation, and plausible causes [28, 59]. Because serious crimes are relatively rare, the sets of evidence involved with such offences tend to be rare also or even unique. Individual features of crime scenarios, and individual types of event and evidence occur significantly more frequently than any specific scenario they may be part of. Variations

on plausible case scenarios may share substantial subsets of situations, events and evidence in common. Therefore, the underlying approach proposed here employs a knowledge representation formalism that models causal relations between typical events and aspects of situations, rather than scenarios as a whole (which conventional case-based approaches tend to make use of). This enables the resulting DSS to be based on a relatively small knowledge base to produce significantly large sets of scenarios.

Nevertheless, scaling up the DSS for generating plausible crime scenarios is constrained by the knowledge acquisition bottleneck, i.e. the difficulty of eliciting and modelling relevant domain knowledge. This has important implications. The completeness of the knowledge base has a significant impact upon the accuracy of the decision support advice that is to be produced. Missing variables and influences may be part of plausible scenarios that explain all the evidence whilst proving a substantially different scenario. Also, missing variables/influences may affect certain conditional probability distributions, though this is a limitation not specific to the present approach but shared by all techniques for knowledge model creation.

The use of first principles, as inspired by model based reasoning, alleviates this concern to a certain extent [23]. However, issues such as the need to consider situations that undermine the *ceteris paribus* assumption make it very difficult to fully exploit the benefits that model-based reasoning offers in the physical systems domain. Therefore, it remains the case that the development of knowledge bases for the system discussed herein requires a substantial amount of effort. The requirement to build a reasonably complete knowledge base prior to the deployment of the DSS is likely to be the greatest obstacle to said deployment.

Model based reasoning with first principles forms the core of this work. Some first principles can be derived from the literature on evidential reasoning. For example, a methodology has been developed to assess witness testimony by separating plausible discrepancies between testimony and hypothetical events in terms of witness veracity, witness objectivity and observational sensitivity of the witness [51]. A Bayesian network has been devised to model two-way trace evidence transfer that is broadly applicable to a wide range of evidence types (not only those that are traditionally considered trace evidence but also, for instance, digital forensics evidence) [2]. Yet, generally speaking, the elicitation of first principles to develop a knowledge base for the system described here is more difficult because it involves the challenging hypothesis synthesis that makes serious crime investigation difficult. Expert practitioners need to be involved with the development of a knowledge base.

Increasingly, expert forensic scientists have been developing Bayesian networks for evidential reasoning in particular domains. For example, most recently, Bayesian networks have been designed to analyse the evidence with respect to hypothetical scenarios which may be encountered in cases of suspected piracy via BitTorrent [37]. Such developments provide formal representations of knowledge that can be ported to a knowledge base required to support the present work, with relative ease. Research on other types of diagram representing evidential reasoning about serious crimes, such as Wigmore charts, also constitutes a potentially valuable source to acquire knowledge from [26].

Given the practical and case-specific nature of forensic science and crime investigation [41], plausible crime scenarios and,

sometimes, the corresponding first principles are only considered during the investigation of a particular case. Tools to represent such scenarios and alternatives between them have recently been developed [3], however. A potentially fruitful avenue of future research would therefore be the development of methods to extract reusable scenario fragments from plausible crime scenarios as they emerge during the course of an investigation. Such techniques would help to automate significant portions of the process of knowledge acquisition, effectively providing a knowledge management approach for forensic laboratories and police forces.

A number of research challenges need to be addressed to support this development, of course. Firstly, representational consistency in the description of hypothetical scenarios and the emerging knowledge base is required. This suggests that a formal ontology would need to be introduced [19]. Secondly, an appropriate mechanism for generalisation of approximate and imprecise first principles is needed. The field of machine learning provides a wealth of methods to potentially accomplish this. However, machine learning approaches typically perform generalisation through speculation that is driven by a method-specific inductive bias [39]. Therefore, and thirdly, a novel technique is necessary to identify emerging inconsistencies in the knowledge base and to correct them during the learning process.

7.6 Usability

The development of a suitable interface to the DSS presented herein is essential to its successful application. As part of the evaluation of the prototypical system that has been implemented (and which is described in [52]), the software has been demonstrated to practitioners, including forensic scientists, police investigators and lawyers. This subsection summarises key findings regarding usability issues [13] that have arisen from this informal, but user-driven part of the evaluation.

The DSS presented here is capable of producing information that police investigators and forensic scientists find particularly useful in a number of aspects. These include: (a) hypothetical explanations of available evidence, (b) additional justifications that discriminate between likely hypotheses (as opposed to evidence that merely confirms preferred hypotheses), (c) representations of plausible scenarios, and (d) rankings of further investigating actions in terms of their potential relative effectiveness. The system is considered to be very helpful, especially for novice investigators. In particular, educational software has been proposed on several occasions, as a potential alternative application (for investigator training).

The information or advice that the present DSS may provide, as discussed in section 6, is not suitable for use in court. The entropy measurements defined therein are based on the posterior probabilities of hypotheses. They vary (ultimately) from the values assigned to the prior probabilities. Prior probabilities of hypotheses can provide a helpful start point in an investigation in the absence of a null hypothesis. It may be based on neutrality between alternative hypotheses (i.e. $P(h_1) = P(h_2) = \dots = P(h_m)$) or reflect expert knowledge (e.g. relative homicide and suicide rates appropriate to the victim's profile). In a court setting, such prior probabilities represent a bias, and prior probabilities are eliminated from the analysis in this setting by employing the likelihood ratio outlined in section 2. However, the Bayesian networks generated by the DSS presented may be employed for this purpose.

Forensic scientists commenting on the usefulness of this approach have pointed out that the causal representation of hypothetical scenarios or larger proportions of the scenario space may be of particular significance. The representation scheme was considered to match well with their scientific approach to an investigation. The prototype software is capable of displaying subsets of a scenario space such as the one shown in Figure 8. The terms in the nodes of the diagram are converted into descriptive text and they can be colour coded to indicate the role (e.g. conjecture, hypothesis, evidence) of each node. One possible view is a single plausible scenario that explains a given set of available evidence or a certain hypothesis. Knowledge produced by the system can thus be presented in a way that is manageable by the user. Another important feature of the causal diagrams is that the nodes can also be colour coded to indicate the relative likelihood of the constituent terms of the diagram, either given a particular scenario (for likelihood ratio calculations) or given a set of observations.

Police investigators generally regard these diagrams as useful but incomplete given the nature of their work. Police officers are particularly constrained by various procedures, policies, perishability of evidence, time and other important considerations such as victim support. For example, each police organisation will have forms, procedures and information systems devised to record and maintain collected evidence, and these are updated as new technologies are introduced (e.g. geotagging). It is crucial that an extended, and hence deployable, version of the prototypical system can be integrated with such existing systems. More generally speaking, the design of an interface that incorporates such information as well for police users remains an important practical issue.

Finally, lawyers tend to employ the information produced by the DSS as part of their arguments. Most of their reasoning typically derives from the evidence. While most practicing lawyers do this with narratives, without the benefit of diagrams, forward-looking legal professionals do employ one of a variety of argument diagramming techniques. Virtually all of these diagramming techniques use evidence rather than hypothetical causes as their origin and reason in a effect-to-cause rather than a cause-to-effect direction [49]. They also employ a richer ontological framework to categorise the role of evidence and premises in an argument (such as so-called primary and ancillary evidence and premises) as well as the way one line of reasoning relates to another. However, the development of such argument structures crucially depends upon the generation of plausible hypotheses, as well as an assessment of the strength of support a given piece of evidence provides for one premise as opposed to another [51]. As such, a valuable line of inquiry is the exploration of the relationships and synergies between Bayesian networks (and scenario spaces in general) on the one hand and argument diagrams on the other hand. Recently, work has been developed to explore such relationships and synergies between Hugin style object-oriented Bayesian networks and Wigmore charts [26].

8 Conclusion

This paper has introduced a novel compositional modelling approach to generating Bayesian networks (BN), called scenario spaces, which describe a range of plausible crime scenarios that are sufficient to explain a given set of available

evidence. A scenario space contains a network of plausible states and events, factually and potentially available evidence, investigating actions, and hypotheses describing important properties of some of the plausible scenarios. It models how these elements of scenarios are causally related and what constraints are imposed over them. Such a space provides an efficient means of storing a wide range of scenarios, because the elements that different scenarios have in common need only be recorded once.

The paper has also extended Bayesian model-based diagnosis techniques to identifying the most likely hypotheses and to refining diagnostic classifications by effective evidence collection strategies. A set of maximum expected entropy reduction techniques have been devised that can identify the investigative actions that are expected to yield the most conclusive evidence. Therefore, the work presented here enables the development of a novel type of decision support system for aiding less experienced crime investigators in constructing effective evidence collection strategies. This has been demonstrated in the paper by means of a realistic application example.

While the proposed approach offers useful functionalities for developing DSSs and shows great potential, a number of further improvements are possible. As the probability distributions captured within the scenario fragments may refer to subjective assessments by experts, of the likely outcomes, which are actually described in vague terms, the use of numeric probabilities conveys an inappropriate degree of precision. Section 7.4 explored a range of extensions to Pascalian probability calculus that can be employed to facilitate elicitation of subjective probabilities. It has indicated why interval-based approaches and their extensions may make it easier for experts to express their beliefs. However, in order to validate this theoretical argument, the use of such probability calculi has to be evaluated either through experiments purposefully designed or by deploying the system for a specific setting and assessing its use. This will constitute a significant piece of empirical research.

Another piece of interesting further research concerns the relaxation of two important assumptions made within this work: (1) probability distributions, which govern the outcomes of different causal influences that may affect the same variable, are independent, and (2) the effects of all causal influences affecting the same variable are combinable using a single type of composition operator. Relaxing these assumptions is essential to allow the approach to be scaled up. It has been demonstrated that the use of qualitative probability calculi and extensions thereof will allow such assumptions to be relaxed through the use of second-order influences, which may affect the existence, strength and direction of other influences [31]. Further development of this work is necessary however, as well as empirical work to evaluate the feasibility of eliciting such knowledge in practice for more complex problems.

As discussed above, the key challenge in future work lies with the acquisition of scenario fragments in general, and the probability distributions therein, in particular. The paper has identified a number of important issues and proposed potentially useful research directions in an effort to address this challenge. However, as the present work has only been evaluated by a relatively small-scale example, further investigation is needed to arrive at a more complete understanding of the nature of the knowledge acquisition bottleneck(s) in this particular problem domain. To accomplish this, a case-study involving a larger scale domain is required, which remains as active research.

Acknowledgements

This work has been supported in part by the Nuffield Foundation Grant NAL/32370 and UK EPSRC grant GR/S63267.

The authors are very grateful to Colin Aitken, Mark Lee, Tom Nelson and Burkhard Schafer for helpful discussions and assistance, whilst taking full responsibility for the views expressed here. The authors are also grateful to the anonymous referees for their constructive comments which have helped improve this paper considerably.

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(V)}
distribution previous-hangings(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}

if {commits-suicide-by(V,M)}
then {suicidal-death(V)}
distribution suicidal-death(V) {
  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,M)}
then {homicidal-death(V)}
distribution homicidal-death(V) {
  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 {accidental-autoerotic-hanging(V)}
then {accidental-death(V)}
distribution accidental-death(V) {
  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(V)}
distribution previous-hangings(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(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}

```

References

- [1] C. Aitken and F. Taroni. *Statistics and the evaluation of evidence*. Wiley, 2004.
- [2] C. Aitken, F. Taroni, and P. Garbolino. A graphical model for the evaluation of cross-transfer evidence in DNA profiles. *Theoretical Population Biology*, 63(3):179–190, 2003.
- [3] F. Bex, S. van den Braak, H. van Oostendorp, H. Prakken, B. Verheij, and G. Vreeswijk. Sense-making software for crime investigation: how to combine stories and arguments? *Law, Probability and Risk*, 6:145–168, 2007.
- [4] A. Biedermann and F. Taroni. Bayesian networks and probabilistic reasoning about scientific evidence when there is a lack of data. *Forensic Science International*, 157(2–3):163–167, 2006.
- [5] M. Bromby, M. MacMillan, and P. McKellar. A CommonKADS representation for a knowledge based system to evaluate eyewitness identification. *International Review of Law Computers and Technology*, 17(1):99–108, 2003.
- [6] H. Chen, J. Schroeder, R.V. Hauck, L. Ridgeway, H. Atabakhsh, H. Gupta, C. Boarman, K. Rasmussen, and A.W. Clements. COPLINK Connect: Information and knowledge management for law enforcement. *Decision Support Systems: Special Issue on Digital Government*, 34(3):271–285, 2002.

- [7] J. Chen and R. Patton. *Robust model-based fault diagnosis for dynamic systems*. Kluwer Academic Publishers, 1999.
- [8] J. Conklin and M. Begeman. gIBIS: a hypertext tool for exploratory policy discussion. *ACM Transactions on Office Information Systems*, 4(6):303–331, 1988.
- [9] 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.
- [10] R. Cook, I. Evett, G. Jackson, P. Jones, and J. Lambert. Case pre-assessment and review in a two-way transfer case. *Science and Justice*, 39:103–111, 1999.
- [11] 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.
- [12] J. de Kleer. An assumption-based TMS. *Artificial Intelligence*, 28:127–162, 1986.
- [13] A. Dix, J. Finlay, G. Abowd, and R. Beale. *Human-Computer Interaction*. Prentice Hall, third edition edition, 2004.
- [14] D. Dixon. Police investigative procedures. In C. Walker and K. Starmer, editors, *Miscarriage of Justice. A Review of Justice in Error*, pages 65–82, 1999.
- [15] M. Druzdzel and M. Henrion. Efficient reasoning in qualitative probabilistic networks. In *Proceedings of the 11th Conference on Artificial Intelligence*, pages 548–553, 1993.
- [16] D. Dzemydiene and V. Rudzkiene. Multiple regression analysis in crime pattern warehouse for decision support. In *Proceedings of the 13th International Conference on Database and Expert Systems Applications*, pages 249–258, 2000.
- [17] I. Evett, G. Jackson, J. Lambert, and S. McCrossan. The impact of the principles of evidence interpretation on the structure and content of statements. *Science and Justice*, 40:233–239, 2000.
- [18] L. Festinger. *A Theory of Cognitive Dissonance*. Row Peterson, 1957.
- [19] D. Gašević, D. Djurić, and V. Devedžić. *Model Driven Engineering and Ontology Development*. Springer, 2nd edition edition, 2009.
- [20] I. Goodman and H. Nguyen. Probability updating using second order probabilities and conditional event algebra. *Information Sciences*, 121(3–4):295–347, 1999.
- [21] 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.
- [22] J. Halliwell and Q. Shen. Linguistic probabilities: theory and applications. *Soft Computing*, 13(2):169–183, 2009.

- [23] W. Hamscher, L. Console, and J. de Kleer, editors. *Readings in Model-Based Diagnosis*. Morgan-Kaufmann, San Mateo, 1992.
- [24] D. Hand, H. Mannila, and P. Smyth. *Principles of Data Mining*. MIT Press, Boston, Massachusetts, 2001.
- [25] R.V. Hauck, H. Atabakhsh, P. Ongvasith, H. Gupta, and H. Chen. Using Coplink to analyze criminal justice data. *IEEE Computer*, 35(3):30–37, 2002.
- [26] A. Hepler, A. Dawid, and V. Leucari. Object-oriented graphical representations of complex patterns of evidence. *Law, Probability and Risk*, 6(1–4), 2007.
- [27] B. Irving and C. Dunningham. Human factors in the quality control of CID investigations and a brief review of relevant police training. *Royal Commission on Criminal Justice Research Studies*, 21, 1993.
- [28] A. Jackson and J. Jackson. *Forensic Science*. Pearsons Education, second edition edition, 2008.
- [29] A. Jamieson. A rational approach to the principles and practice of crime scene investigation: I. principles. *Science & Justice*, 44(1):3–7, 2004.
- [30] S. Kaza, Y. Wang, and H. Chen. Suspect vehicle identification for border safety with modified mutual information. In *Proceedings of the IEEE International Conference on Intelligence and Security Informatics*, pages 308–318, 2006.
- [31] J. Keppens. Towards qualitative approaches to bayesian evidential reasoning. In *Proceedings of the 11th International Conference on Artificial Intelligence and Law*, pages 17–25, 2007.
- [32] J. Keppens and B. Schafer. Knowledge based crime scenario modelling. *Expert Systems with Applications*, 30(2):2003–222, 2006.
- [33] 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.
- [34] 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.
- [35] J. Keppens, Q. Shen, and B. Schafer. Probabilistic abductive computation of evidence collection strategies in crime investigation. In *Proceedings of the 10th International Conference on Artificial Intelligence and Law*, pages 215–224, 2005.
- [36] J. Keppens and J. Zeleznikow. A model based reasoning approach for generating plausible crime scenarios from evidence. In *Proceedings of the 9th International Conference on Artificial Intelligence and Law*, pages 51–59, 2003.
- [37] M. Kwan, K. Chow, F. Law, and P. Lai. Reasoning about evidence using bayesian networks. In *Advances in Digital Forensics IV*, International Federation for Information Processing, pages 275–289. Springer, 2008.

- [38] R. Loui, J. Norman, J. Alpeter, D. Pinkard, D. Craven, J. Lindsay, and M. Foltz. Progress on room 5: a testbed for public interactive semi-formal legal argumentation. In *Proceedings of the 6th International Conference on Artificial Intelligence and Law*, pages 207–214, 1997.
- [39] T.M. Mitchell. *Machine Learning*. McGraw-Hill, 1997.
- [40] J. Mortera, A. Dawid, and S. Lauritzen. Probabilistic expert systems for dna mixture profiling. *Theoretical Population Biology*, 63(3):191–205, 2003.
- [41] J. Nordby. Here we stand: What a forensic scientist does. In S. James and J. Nordby, editors, *Forensic Science: An Introduction to Scientific and Investigative Techniques*, pages 1–14. CRC Press, 2005.
- [42] G. Oatley, B. Ewart, and J. Zeleznikow. Decision support systems for police: lessons from the application of data mining techniques to "soft" forensic evidence. *Artificial Intelligence and Law*, 14(1):35–100, 2006.
- [43] S. Parsons. Refining reasoning in qualitative probabilistic networks. In *Proceedings of the 11th Conference on Uncertainty in Artificial Intelligence*, 1995.
- [44] S. Parsons. Qualitative probability and order of magnitude reasoning. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, 11(3):373–390, 2003.
- [45] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan-Kaufmann, 1988.
- [46] D. Poole. Explanation and prediction: An architecture for default and abductive reasoning. *Computational Intelligence*, 5(2):97–110, 1989.
- [47] H. Prakken, C. Reed, and D. Walton. Argumentation schemes and generalisations in reasoning about evidence. In *Proceedings of the 9th International Conference on Artificial Intelligence and Law*, pages 215–224, 2003.
- [48] M. Redmond and C. Blackburn Line. Empirical analysis of case-based reasoning and other prediction methods in a social science domain: Repeat criminal victimization. In *Proceedings of the 5th International Conference on Case-Based Reasoning*, pages 452–464, 2003.
- [49] C. Reed, D. Walton, and F. Macagno. Argument diagramming in logic, law and artificial intelligence. *Knowledge Engineering Review*, 22(1):87–109, 2007.
- [50] S. Renooij, L. van der Gaag, and S. Parsons. Context-specific sign-propagation in qualitative probabilistic networks. *Artificial Intelligence*, 140:207–230, 2002.
- [51] D. Schum. *The Evidential Foundations of Probabilistic Reasoning*. Northwestern University Press, 1994.
- [52] Q. Shen, J. Keppens, C. Aitken, B. Schafer, and M. Lee. A scenario-driven decision support system for serious crime investigation. *Law, Probability and Risk*, 5(2):87–117, 2006.

- [53] F. Taroni, C. Aitken, P. Garbolino, and A. Biedermann. *Bayesian networks and probabilistic inference in forensic science*. John Wiley and Sons Ltd., Chichester, 2006.
- [54] J. Toland and B. Rees. Applying case-based reasoning to law enforcement. *International Association of Law Enforcement Intelligence Analysts Journal*, 15(1):106–125, 2005.
- [55] B. Verheij. Automated argumentation assistance for lawyers. In *Proceedings of the 7th International Conference on Artificial Intelligence and Law*, pages 43–52, 1999.
- [56] J. von Plato. de finetti’s earliest works on the foundations of probability. *Erkenntnis*, 31:263–282, 1989.
- [57] P. Walley. Statistical inferences based on a second-order possibility distribution. *International Journal of General Systems*, 9:337–383, 1997.
- [58] D. Walton. Argumentation and theory of evidence. *New Trends in Criminal Investigation and Evidence*, 2:711–732, 2000.
- [59] P. White. *Crime Scene to Court: The Essentials of Forensic Science*. The Royal Society of Chemistry, second edition edition, 2004.
- [60] Y. Xiang, M. Chau, H. Atabakhsh, and H. Chen. Visualizing criminal relationships: comparison of a hyperbolic tree and a hierarchical list. *Decision Support Systems*, 41:69–83, 2005.