Probabilistic Abductive Computation of Evidence Collection Strategies in Crime Investigation

Jeroen Keppens
Dept. of Computer Science
University of Wales
Penglais Campus
Aberystwyth, SY23 3DB, UK
jrk@aber.ac.uk

Qiang Shen
Dept. of Computer Science
University of Wales
Penglais Campus
Aberystwyth, SY23 3DB, UK
qqs@aber.ac.uk

Burkhard Schafer
School of Law
The University of Edinburgh
Old College, South Bridge
Edinburgh EH8 9YL, UK
bschafer@staffmail.ed.ac.uk

ABSTRACT

This paper presents a methodology for integrating two approaches to building decision support systems (DSS) for crime investigation: symbolic crime scenario abduction [16] and Bayesian forensic evidence evaluation [5]. This is achieved by means of a novel compositional modelling technique that allows for automatically generating a space of models describing plausible crime scenarios from given evidence and formally represented domain knowledge. The main benefit of this integration is that the resulting DSS is capable to formulate effective evidence collection strategies useful for differentiating competing crime scenarios. A running example is used to demonstrate the theoretical developments.

Keywords

Model Based Reasoning, Bayesian Networks, Decision Support Systems, Crime Investigation

1. INTRODUCTION

Evidence is the central consideration in any proper crime investigation effort. It has been the subject of a considerable amount of research aimed at developing software systems to support human decision making. Typical examples of such work include establishing the validity of evidence [2], identifying plausible crimes on the basis of evidence [3, 13, 24, 26], formulating hypotheses for individual cases on the basis of evidence and evaluating their likelihood [1, 4, 20], and establishing appropriate legal arguments in court proceedings on the basis of evidence [18, 23, 27, 29]. However, relatively little work has focussed on developing mechanisms useful to identify effective evidence collection strategies. Yet, this is an important issue as failures to consider crucial lines of inquiry are a prominent cause of miscarriages of justice [8]. This paper introduces an approach to tackle this problem.

Constructing evidence collection strategies is a relatively complex problem as it requires knowledge of how evidence was/will be formed and a means of evaluating the likelihood of obtaining evidence under hypothetical scenarios [16]. The work presented herein addresses this by integrating two novel ideas that have recently emerged from research in building decision support systems (DSS) for crime investigation. The first idea is that a space of plausible scenarios that each explains the available evidence in a case can be efficiently stored in an assumption based truth maintenance system (ATMS) [16]. The second is that the Bayesian inference method forms an effective means for evaluating how well a given piece of evidence can adjust belief in one scenario over a possible alternative [5].

This paper will show how a compositional modelling method can extend a symbolic scenario space (within the ATMS paradigm) with a Bayesian representation of it, and how the resulting extended scenario can be analysed to formulate effective evidence collection strategies. It is organised as follows: Section 2 summarises the underlying work upon which this paper is based. Section 3 presents an outline of the proposed work. Then, Section 4 shows how Bayesian Networks (BNs) can be synthesised from knowledge, and Section 5 describes how such BNs can be analysed to compute effective evidence collection strategies. Finally, Section 6 concludes this paper and presents some interesting areas of future work.

2. BACKGROUND

2.1 Compositional modelling method

The present work extends the compositional modelling approach previously introduced in [16], which generates symbolic descriptions of crime scenario from evidence. The approach employs a knowledge base that consists primarily of so-called scenario fragments. Scenario fragments are production rules, of the form `if` *conjunction of antecedent predicates* `assuming` *conjunction of assumption predicates* `then` *consequent predicate*, which describe certain causal relationships. Here, assumption predicates are pieces of uncertain information whose truth must be established by considering their consequences. For example,

```
if {suffers(P,C),
    cause-of-death(P,C),
    medical-examiner(E)
} assuming {
    determine(E,cause-of-death(P)),
    correct-diagnosis(E,cause-of-death(P))
} then {
    cause-of-death-report(E,P,C)}
```

states that if a person `P` suffers from an ailment or injury `C`, `C` is the cause of death of `P`, and there is a medical examiner `E`, *and*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
ICAIL '05, June 6-11, 2005, Bologna, Italy.
Copyright 2005 ACM 1-59593-081-7/05/0006 ...$5.00.

assuming that E determines the cause of death of P and entails a correct diagnosis, then there will be a piece of evidence in the form of a cause of death report indicating that according to E, the cause of death of P is C.

The knowledge base also contains so-called inconsistencies, or rules describing which combinations of states and events are contradictory with one another, and therefore can not be true within a single scenario. For instance,

```
inconsistent {
    suicide(P,hanging),
    is-hanged(P)}
```

states that a person P can not both commit suicide by hanging and be hanged by someone else. This knowledge base, of scenario fragments and inconsistencies is employed to create a scenario space. A scenario space is a hypergraph representation in which the nodes correspond to plausible states and events, and hyperarcs from a set of antecedent nodes to a single consequent node describe a causal relation from the conjunction of the antecedents to the consequent. The scenario space is constructed by a computational procedure that applies the scenario fragments first in a backward chaining fashion to instantiate all plausible scenarios that may explain the available evidence and then, in a forward chaining fashion to instantiate any possible evidence that can be collectible under certain generated plausible scenarios.

Because a realistic scenario space would be too large to show in this paper, Figure 1 depicts a sample part of the structure of a scenario space. This sample space describes different plausible scenarios of a dead body found hanging as well as the possible additional evidence collectible under certain scenarios. The plausible scenarios include suicide by hanging, a number of potential homicide situation and non-suicide cases where the victim engaged in autoerotic activities involving hanging and died accidentally. The different squares in the figure correspond to a set of nodes and hyperarcs between them, showing part of one or more scenarios (the reader is referred to [16] for a more detailed example of a scenario space).

The nodes in the scenario space represent plausible states and events that are part of one or more possible scenarios. There are also special types of node that convey additional information that may aid in decision support. These concepts have been adapted from earlier work on (probabilistic) abductive reasoning [22, 25] and model based diagnosis [12]. In particular, *evidence* nodes are pieces of known information that are deemed to be observable consequences of a possible crime[1]. *Facts* are pieces of known information that do no require an explanation. In practice, it is often convenient to accept some information at face value without elaborating possible justifications. For instance, when a person is charged with analysing the handwriting on the aforementioned suicide node, 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 nodes 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 node 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.

In the knowledge base, 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 evidence sets, in the case of investigative actions). Conjectures contained in the knowledge base are identified in the `assuming` clause of the scenario fragments.

The scenario space is generated and accessed through an assumption based truth maintenance system (ATMS) [7]. Given one or more nodes in the scenario space, the ATMS can efficiently determine all the minimal sets of assumptions that logically (and in the present application causally) entail the given nodes, and that are not inconsistent. This functionality of the ATMS is employed to extract the following information, which may be helpful for crime investigators (for a more detailed explanation of this approach, the reader is again referred to [16]):

- *Which scenarios are plausible?* Scenarios in the scenario space are associated with consistent sets of assumptions that logically entail the available evidence. Therefore, any scenario containing a set of assumptions W and all the nodes that follow from W in the scenario space is plausible provided W entails all the available evidence.

- *Is a given hypothetical predicate supported by the available evidence?* A hypothesis is supported by the available evidence if there is a plausible scenario that entails it.

- *What evidence may be available under a given scenario or hypothesis?* A piece of evidence is available under a given scenario or hypothesis if a consistent collection of assumptions exist that jointly entail both the scenario or hypothesis and the piece of evidence.

2.2 Bayesian approach to evidence evaluation

Briefly, the methodology of the Bayesian approach to evaluating a piece of forensic evidence e, proposed by Evett et. al. [9], 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_{defense}. 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 $P(e \mid p_{\text{defense}})$ of obtaining the given piece of evidence in the defence scenario. One approach to modelling these probabilities is to

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

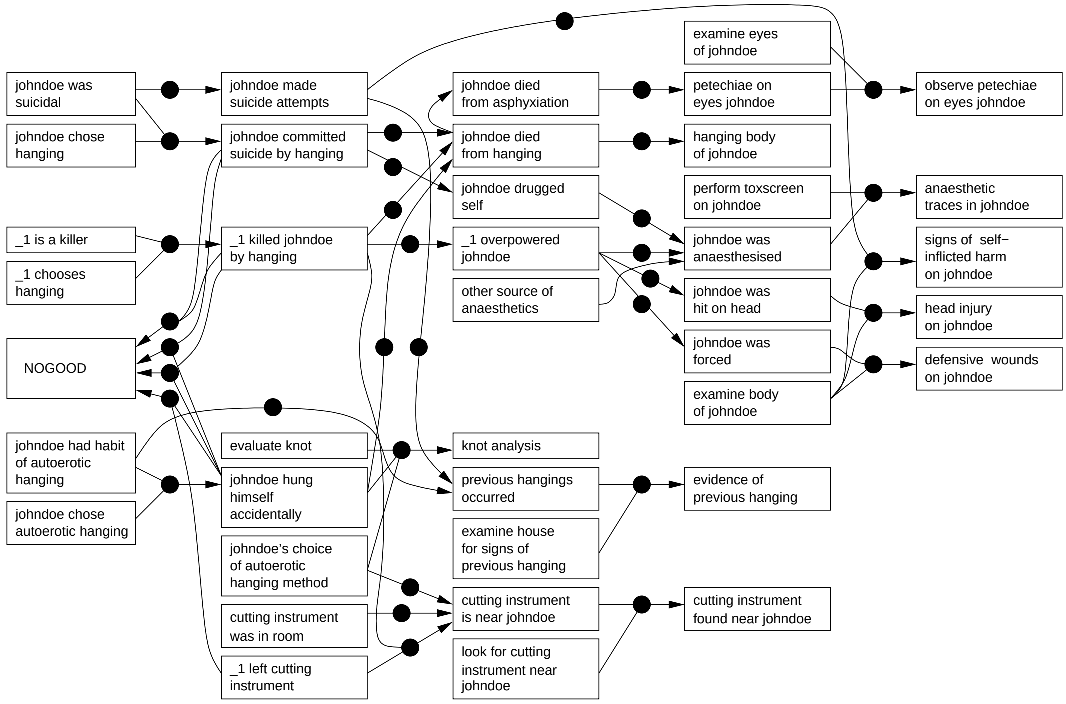

Figure 1: Sample scenario subspace structure

use Bayesian networks. Bayesian networks 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*:

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

The greater LR is (above 1), 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 lab 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 [14]. The specific use of precise probabilities is more controversial, although it is adopted by major forensic laboratories, such as the UK's Forensic Science Service [4]. 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 [20]. Nevertheless, the approach can also be successfully applied to cases where analyst has to rely on subjective probabilities, by performing a detailed sensitivity analysis [1] or by incorporating uncertainty concerning the probability estimates within the Bayesian model [11].

3. ARCHITECTURE

Figure 2 shows the architecture of a decision support system (DSS) that reflects the proposed approach taken in this work. The DSS essentially works in two phases. In the first phase, the *synthesis* methods generate a space of plausible scenarios. This will result in two versions of the scenario space: 1) a symbolic version stored in an assumption based truth maintenance system (ATMS), from which minimal consistent scenarios can be derived, and 2) a numerical version expressed by a Bayesian network (BN) that contains the probabilistic information. In the second phase, the scenario space is *analysed* to generate decision support information. As shown in [16], the ATMS contains information sufficient to answer the three types of query as described in Section 2.1. The rest of this paper will show how the combination of the ATMS and the BN enables the construction of evidence collection strategies. Naturally, Figure 2 also includes a feedback loop indicating that investigators may use the decision support information to collect additional evidence.

4. SCENARIO SPACE SYNTHESIS

This section describes the synthesis method for creating Bayesian scenario spaces, with Section 4.1 addressing the knowledge representation and Section 4.2 the synthesis algorithm.

4.1 Knowledge base

4.1.1 Concepts

While the symbolic approach introduced in [16] aims at aiding the user/investigator to establish the truth or falsehood of hypothetical scenarios, states and events that are represented by means of predicates, a probabilistic reasoner aims to compute probability distributions for variables (describing, say, the intensity of contact between a murderer and his victim or the amount of trace material transferred between two people) with a given domain. To facilitate the integration of both techniques so as to maximise their potential benefits, the subject of the reasoner proposed herein are tuples $\langle p, D_p, v_p, \oplus \rangle$. Each such tuple corresponds to a variable,

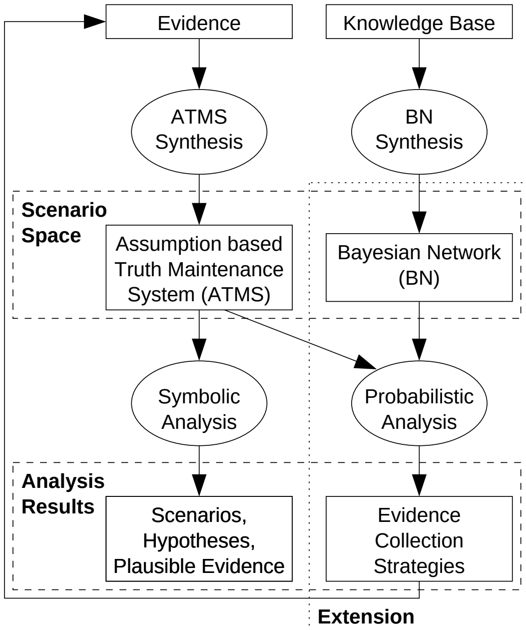

Figure 2: Extended architecture

which is identified by a *predicate* p, which has a *domain* D_p of values, including a default value $v_p \in D_p$, that can be assigned to the variable, and which is associated with a *combination operator* $\oplus : D_p \times D_p \mapsto D_p$ that describes how the effects of different influences acting upon the variable are combined.

As in previous work [16], scenario fragments describe causal influences among types of state and event and crime scenarios. However, the consequence of any influence is no longer presumed to be deterministic, but governed by predefined probability distributions. Thus, the notion of scenario fragment is extended to incorporate a set of probability distributions, one for each combination of the antecedent and assumption variables. As such, probabilistic scenario fragments are represented by:

if $\{p_1, \ldots, p_k\}$
assuming $\{p_l, \ldots, p_m\}$
then $\{p_n\}$
distribution p_n $\{$

$\vdots$

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

$\vdots$ $\}$

where $\{p_1, \ldots, p_k\}$ is the set of antecedent predicates, $\{p_l, \ldots, p_m\}$ is the set of assumption predicates, p_n is the consequent predicate, each v_i is a value taken from the domain D_{p_i} of the variable identified by p_i and each q_j is a real value in the range $[0, 1]$. Similar to conventional scenario fragments [16], the `if`, `assuming` and `then` components of probabilistic scenario fragments respectively describe the types of antecedents, assumptions and consequent of a causal relation. As opposed to conventional scenario fragments, however, the antecedents, assumptions and consequent within a fragment refer to variables instead of truth assignments. The domain value that a consequent has is influenced by the values of the antecedent and assumption variables and the probability distributions defined in the `distribution` component of the fragment. In particular, each line

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

defines a discrete probability distribution

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

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

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

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

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

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

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

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

if $\{p_1, \ldots, p_k\}$
then $\{$nogood$\}$
distribution nogood $\{$
$v_1, \ldots, v_k$ -> $\top : 1, \ldots, \bot : 0\}$

(5)

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. Therefore, $P(p_1 : v_1, \ldots, p_k : v_k \mid \texttt{nogood} : \bot) = 0$.

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 $\{v_1 : q_1, \ldots, v_j : q_j\}$

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

4.1.2 Presumptions

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

1. *Any two probability distributions taken from two scenario fragments involving the same consequent variable are independent*. Intuitively, this assumption indicates that the outcome of an influence implied by one scenario fragment is not affected by that of another.

2. *There are no cycles in the knowledge base*. This means that there is no subset of scenario fragments in the knowledge base which is of the form $\langle\{\ldots,p_1,\ldots\},P_{a1},p_2,f_1\rangle$, $\langle\{\ldots,p_2,\ldots\},P_{a2},p_3,f_2\rangle,\ldots,\langle\{\ldots,p_n,\ldots\},P_{an},p_1,f_n\rangle$. This assumption is required because BNs can not represent such information as they are inherently acyclic [21].

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 4.2.2). When only a few different scenario fragments are needed to describe phenomena that affect the same variable, it may be possible to satisfy the independence presumption by modelling the correlation between the phenomena by means of latent variables [17]. In larger application domains, where there may be many scenario fragments affecting a single variable, this is no longer practical. Future work will seek to relax this assumption in order to generalise further the application of the method proposed.

4.2 Algorithm

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

4.2.1 Structure

The procedure to synthesise the structure of a Bayesian scenario space is an extension of the symbolic scenario space construction algorithm introduced in [16]. First, the algorithm composes a hypergraph representation of the scenario space by means of the same procedure as given in [16]. This results in a hypergraph $\langle N,A,J\rangle$, where N is a set of nodes, A is a subset of N containing all the assumption nodes in the hypergraph and J maps each node $n_i \in N$ to a set $J(n_i)$ of justification sets. The latter is best explained by an example. For instance, let $n_1,\ldots,n_5$ represent the following states and events in the scenario space of Figure 1:

n_1 johndoe committed suicide by hanging
n_2 unknown person _1 killed johndoe by hanging
n_3 johndoe hung himself accidentally
n_4 johndoe's choice of autoerotic hanging method
n_5 knot

Then, $J(n_5)=\{\{n_1\},\{n_2\},\{n_3,n_4\}\}$. Inconsistencies are processed in the same manner, but they all have the same nogood node as their consequent, which will be identified by $n_\perp$ in what follows.

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

$$P(p:x)=\begin{cases}1 & \text{if } x=v\\ 0 & \text{otherwise}\end{cases} \tag{7}$$

where $x\in D_p$. Therefore, each root node in the hypergraph generated by this procedure must be either a fact or an assumption. However, the backward chaining phase can not guarantee this. Hence, when the procedure terminate, all those nodes, which were originally regarded as root nodes and which are not a fact or assumption and all the justifications including these nodes are deemed spurious. Consider, for instance, the following scenario fragment:

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

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

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

While spurious nodes are ignored in an ATMS, if the belief propagation algorithm of a BN still attempts to take them into account, this would inevitably result in an error as spurious nodes have no prior probability distribution. Therefore, spurious nodes and justifications must both be removed from the hypergraph. This work employs the following procedure to accomplish this in a given hypergraph $\langle N,A,J\rangle$:

Algorithm 4.1: SPURIOUSNODEREMOVAL($\langle N,A,J\rangle,\mathbf{K}$)

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

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

Algorithm 4.2: CREATEDAG($\langle N,A,J\rangle$)

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

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

4.2.2 *Conditional Probability Tables*

A BN also requires a complete specification of the conditional probability tables to be of any practical use. Let m be the number of states of each node in the BN and q be the number of parents of each non-root node. Then, a total of $m^q \times (m-1)$ probabilities must be assigned to each non-root node. In an abductive diagnosis application, q may become large, thereby inhibiting the manual specification of the conditional probabilities. For example, the probability distribution of the amount of a particular anaesthetic in the blood of a victim's body can be affected by self-medication, consumption of a spiked drink, surgery, etc.

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

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

$$P(p_n : c \mid A) = P\left[\bigvee_{c_1 \oplus \ldots \oplus c_k = c} \left(\bigwedge_{i=1,\ldots,k} (A \stackrel{C_i}{\rightarrow} p_n : c_i)\right)\right] \quad (8)$$

According to (8), computing $P(p_n : c \mid A)$ involves calculating the likelihood of a combination of events described by a disjunctive normal form (DNF) expression. Because the occurrence of different combinations of outcomes $c_1, \ldots, c_k$ of the scenario fragments $C_1, \ldots, C_k$ involves mutually exclusive events, the calculation can be resolved by adding the probabilities of the conjuncts in (8):

$$P(p_n : c \mid A) = \sum_{c_1 \oplus \ldots \oplus c_k = c} P\left(\bigwedge_{i=1,\ldots,k} (A \stackrel{C_i}{\rightarrow} p_n : c_i)\right) \quad (9)$$

From presumption 1, the outcomes of different scenario fragments (with the same consequent), in case of a given set of assignments of the antecedent and assumption variables, correspond to independent events. Therefore, the probability of the conjunctions in (9) is equal to the product of the probabilities of their conjuncts, and (9) is calculated as follows:

$$P(p_n : c \mid A) = \sum_{c_1 \oplus \ldots \oplus c_k = c} \left(\prod_{i=1,\ldots,k} P(A \stackrel{C_i}{\rightarrow} p_n : c_i)\right) \quad (10)$$

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

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

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

where `autoerotic-hanging-habit(V)` and `previous-suicide-attempts(V)` correspond to boolean variables, and `previous-hanging(V)` to a variable taking values from the domain {`never`, `veryfew`, `several`} defined over combination operator $\max$[2]. Then, the probabilities of assignments to `previous-hanging(V)`, given that `autoerotic-hanging-habit(V)` and `previous-suicide-attempts(V)` are assigned $\top$, can be computed as follows:

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

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

Thus, $P(p_3 : \texttt{veryfew} | p_1 : \top, p_2 : \top)$ can be computed as follows, according to (10):

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

Similarly, it can be shown that

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

5. SCENARIO SPACE ANALYSIS

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

[2] Thus `never`$\oplus$`veryfew` = `veryfew` and `veryfew`$\oplus$`several` = `several`

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

The work will be illustrated by means of probabilities derived from a BN which has been generated using the techniques of Section 4 and which is a Bayesian representation of the symbolic scenario space given in Figure 1. The probabilities themselves are arbitrarily chosen and merely serve as a means to illustrate the calculations. Also, due to space limitations, the knowledge base and the BN can not be provided in this paper but interested readers can find it online at http://users.aber.ac.uk/jrk/examples/icail2005.html.

5.1 Hypothesis sets and query types

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

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

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

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

are respectively associated with the following hypothesis sets:

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

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

```
if {p_1, ..., p_n}
then {p_{n+1}}
distribution p_{n+1} {⊥, ..., ⊥->⊤ : 1, ⊥ : 0}
```

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

```
inconsistent {p_i : ⊤, p_j : ⊤}
```

5.2 Entropy

The work here employs an information theory based approach, which is widely used in areas such as machine learning [19] and model based diagnosis [12]. Information theory utilises a measurement of doubtness over a range of choices, called entropy. Applied to the present problem, the *entropy* over an exhaustive set of mutually exclusive hypotheses $H = \{h_1, \ldots, h_m\}$ is given by:

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

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

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

$$\epsilon(H \mid E) = -\sum_{h \in H} P(h \mid E) \log P(h \mid E) \tag{13}$$

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

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

Thus, as an instance,

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

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

$$\min_{a \in A} E(\epsilon(H \mid E), a) \tag{15}$$

Note that the entropy values calculated by equation (13) are affected by the prior distributions assigned to assumptions, as described in 4.1.1. Within the context of evidence evaluation (which is the conventional application of the likelihood ratio approach), this is a controversial issue as decision regarding the likelihood of priors, such as the probability that a victim had autoerotic hanging habits, are a matter for the courts to decide on. In the context of an investigation, however, these prior distributions may provide helpful information often ignored by less experienced investigators. For example, the probability of suicides or autoerotic deaths are often underestimated. As such, decision criterion (15) is a useful means of deciding on what evidence to collect next. Yet, the minimal entropy decision rule does not yield information that should be used for evidence evaluation in court.

5.3 Minimal entropy-based evidence collection

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

$$E(\epsilon(H \mid E), a) = \sum_{e \in E_a} P(e \mid a : \top, E)\epsilon(H \mid E \cup \{a : \top, e\}) \quad (16)$$

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

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

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

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

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

5.4 Extensions

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

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

$$\begin{aligned}
&E(\epsilon(H \mid E), a_1, \ldots, a_v) \\
&\quad = \sum_{e_1 \in E_{a_1}} \cdots \sum_{e_v \in E_{a_v}} P(e_1, \ldots, e_v \mid a_1 : \top, \ldots, a_v : \top, E) \quad (18) \\
&\qquad\qquad \epsilon(H \mid e_1, a_1 : \top, \ldots, e_v, a_v : \top, E)
\end{aligned}$$

In order to determine $E(\epsilon(H \mid E), a_1, \ldots, a_v)$, equation (18) can be simplified as follows:

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

5.4.2 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 resulting symptoms may occur in any combination of such pieces of evidence. The above approach can be readily extended to account for this by computing the EPEs after performing action a with associated evidence sets $E_{a,1}, \ldots, E_{a,w}$:

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

5.4.3 *Multiple hypothesis sets*

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

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

5.5 User interface

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

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

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

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

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

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

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

Table 1: Evaluation of investigative actions

6. CONCLUSIONS AND FUTURE WORK

This paper has introduced a novel decision support system aimed at aiding crime investigators in establishing appropriate evidence collection strategies. It extends previous work [16], in an effort to generate a space of models of plausible scenarios that is useful to explain what could have caused the available evidence in a given case, with probabilistic information. Firstly, conventional scenario fragments have been augmented to represent causal influences with a non-deterministic outcome governed by probability distributions. Secondly, a new scenario space generation algorithm is proposed that employs these probabilistic scenario fragments to create both the structure and the conditional probability tables of Bayesian Network representation of the scenario space. And finally, a minimal expected entropy technique has been devised to analyse the Bayesian Scenario Space for producing effective evidence collection strategies.

The approach described herein has been integrated into a prototype decision support system that enables users to access this functionality and to visualise generated scenarios and inferred decision support information (although the actual presentation of this system is left out due to space limitations). With the ongoing development of larger knowledge bases, it is hoped that this approach will lead to a novel type of decision support system to guide less experience crime investigators in considering appropriately broad hypothesis sets, formulating effective evidence collection strategies and henceforth avoiding miscarriages of justice.

While the proposed approach presented herein offers very useful functionalities for DSS, a number of further improvements are possible. As the probability distributions in the scenario fragments refer to subjective assessments by experts of the likely outcomes, which are described in terms of vague concepts, the use of numeric probabilities conveys an inappropriate degree of precision. It would be more appropriate to incorporate a measurement of imprecision within the probability distributions. A number of approaches can provide a means of representing and reasoning with such imprecision, such as second-order probability theory [6, 10, 28] and linguistic probability theory [11]. Investigation into the use of symbolic probabilities forms an interesting immediate future work.

In future work, the prototype application will be extended to include more detailed and sophisticated scenarios. This requires the development of a larger scale common-sense reasoning knowledge base, which may constitute a substantial challenge. One consideration is that typical applications of this work involve reasoning about hypothetical scenarios that occur in time and space. The likelihood of such scenarios is not only affected by the observed symptoms or evidence, but also by constraints on the time and space in which the events in the scenarios occur. Therefore, further research into incorporating temporal and spatial reasoning in this framework is of significant relevance to this work.

Other important future work concerns relaxing two important assumptions made within this work: 1) probability distributions governing the outcomes of different causal influences (and hence represented in distinct scenario fragments) that affect the same variable must be independent, and 2) the effects of all causal influences affecting the same variable must be combinable using a single composition operator. It has been argued in this paper that these issues can be overcome by adding appropriate variables to the scenario fragments in question and that the inconvenience posed by these additional variables is far outweighed by the benefits of compositionality of scenario fragments. However, the knowledge representation scheme adopted seems to allow the aforementioned assumptions to be relaxed. For example, information on the correlation between causal influences, specified by scenario fragments, could be added to the knowledge base, thereby explicitly representing how influences are interdependent. Yet, exactly how this may be implemented requires considerable further studies. Also, multiple composition operators can be allowed by defining rules of composition, as in the work on compositional model repositories [15].

Acknowledgements

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

7. REFERENCES

[1] 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.

[2] M. Bromby, M. MacMillan, and P. McKellar. A common-kads representation for a knowledge based system to evaluate eyewitness identification. *International Review of Law Computers and Technology*, 17(1):99–108, 2003.

[3] 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.

[4] 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.

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

[6] 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.

[7] J. de Kleer. An assumption-based TMS. *Artificial Intelligence*, 28:127–162, 1986.

[8] D. Dixon. Police investigative procedures. In C. Walker and K. Starmer, editors, *Miscarriage of Justice. A Review of Justice in Error*, 1999.

[9] 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.

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

[11] 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.

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

[13] 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.

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

[15] 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.

[16] 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.

[17] P. Lazarsfeld and N. Henry. *Latent Structure Analysis*. Houghton Mill, Boston, 1968.

[18] R. Loui, J. Norman, J. Alpeter, D. Pinkard, D. Craven, J. Linsday, and M. Foltz. Progress on room 5: a testbed for public interactive semi-formal legal argumentation. In *Proceedings of the 6th Internation Conference on Artificial Intelligence and Law*, pages 207–214, 1997.

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

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

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

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

[23] H. Prakken, C. Reed, and D. Walton. Argumentation schemes and generalisations in reasoning about evidence. In *Proceedings of the 9th Internation Conference on Artificial Intelligence and Law*, 2003.

[24] 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.

[25] D. Schum. *The Evidential Foundations of Probabilistic Reasoning*. Northwestern University Press, 1994.

[26] J. Toland and B. Rees. Applying case-based reasoning to law enforcement. *International Association of Law Enforcement Intelligence Analysts Journal*, 15, 2005.

[27] B. Verheij. Automated argumentation assistance for lawyers. In *Proceedings of the 7th Internation Conference on Artificial Intelligence and Law*, pages 43–52, 1999.

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

[29] D. Walton. Argumentation and theory of evidence. *New Trends in Criminal Investigation and Evidence*, 2:711–732, 2000.