Aberystwyth University

A scenario-driven decision support system for serious crime investigation
Shen, Qiang; Keppens, Jeroen; Aitken, Colin; Schafer, Burkhard; Lee, Mark

A Scenario Driven Decision Support System

for Serious Crime Investigation

Qiang Shen,* Jeroen Keppens,† Colin Aitken,‡ Burkhard Schafer,§ and Mark Lee ¶

Abstract

Consideration of a wide range of plausible crime scenarios during any crime investigation is important to seek convincing evidence and hence to minimise the likelihood of miscarriages of justice. It is equally important for crime investigators to be able to employ effective and efficient evidence collection strategies that are likely to produce the most conclusive information under limited available resources. An intelligent decision support system that can assist human investigators by automatically constructing plausible scenarios and reasoning with the likely best investigating actions will clearly be very helpful in addressing these challenging problems. This paper presents a system for creating scenario spaces from given evidence, based on an integrated application of techniques for compositional modelling and Bayesian network-based evidence evaluation. Methods of analysis are also provided by the use of entropy to exploit the synthesised scenario spaces in order to prioritise investigating actions and hypotheses. These theoretical developments are illustrated by realistic examples of serious crime investigation.

Keywords: Crime investigation; Decision support; Scenario generation; Scenario fragments; Bayesian networks; Evidence evaluation; Conditional independence; System architecture; Entropy.

1 Introduction

In the criminal justice system, practitioners from highly diverse backgrounds and professional cultures, including police, crime scene examiners, scientists, medical examiners, legal professionals, policy makers and researchers, have to interact on a recurrent basis. As a result, crime investigations take place at the intersection of numerous and often conflicting demands by the various stake-holders. The heterogeneity of these multi-disciplinary groups poses serious challenges. Some of the core issues involved in the crime scene to court process are the understanding, integration and communication of the different spheres of knowledge, concepts and experience – concepts

*Department of Computer Science, University of Wales, Aberystwyth, SY23 3DB, UK

†Department of Computer Science, King's College London, London WC2R 2LS, UK

‡School of Mathematics, University of Edinburgh, Edinburgh EH9 3JZ, UK

§School of Law, University of Edinburgh, Edinburgh EH8 9YL, UK

¶Department of Computer Science, University of Wales, Aberystwyth, SY23 3DB, UK

and experiences which will not always match. For a scientist, two blood samples either match or do not match, regardless of whether the samples were obtained with or without the appropriate warrant. The need for the police to stay within budgetary constraints in their work is only partly recognised in the law of evidence. Computer technology has the potential of facilitating cross disciplinary communication, first by making underlying assumptions and procedures explicit, and second by facilitating access to the relevant expert knowledge by the non-expert.

Previous investigations into the analysis of high profile miscarriages of justice in the 1980s, as reported in numerous papers (e.g. [7, 14, 26]), have led to an interesting finding: Many of these cases may be understood as a clash between professional cultures and their respective reward structures. Scientists, to some extent at least, adhere to a falsificationist model of inquiry. In this approach, conflicting explanations of the available evidence are actively sought and then eliminated in a process of rigorous testing. Police investigators by contrast, tend to formulate a case theory at a very early stage of an investigation, and then try inductively to confirm that theory to an extent high enough to convince the prosecution service to indict the suspect. This can lead to a situation where exculpating evidence is overlooked [10].

To help address this problem, a computer-aided tool that makes use of human knowledge for crime investigation has been built, which ensures as complete a space of likely cases as possible to be considered [22]. This work was inspired by solutions developed for the diagnosis of complex physical systems [12]. It involves two conceptually distinct stages. In the first, *constructive* stage, the computer system develops abductively, on the basis of the established facts, a range of alternative explanations to the case theory of the investigator. All of these should account equally for the available evidence. The aim is to keep track of alternative explanations of the evidence and to remind the investigating officer that not all avenues of inquiry have as yet as been exhausted. In the second, *critical* stage, the system suggests which additional tests and inquiries can be carried out to decide between the different explanations. For this, it deduces what other facts would be expected to be found if the respective theories were true. Of particular interest are those facts that are inconsistent with an investigating hypothesis, thereby helping to eliminate provably wrong explanations.

The most intuitive illustration of this idea is, perhaps, through the exploitation of the concept of *alibi*. Suppose that the evidence which triggers the investigation is the discovery of the body of Smith, with a knife in his back. The preferred police hypothesis is that Jones stabbed Smith. While this would explain how the evidence was created, everything found so far may still be consistent with the guilt of an alternative suspect, Miller. In the first stage, the system builds these two alternative explanations for the evidence discovered. From the hypothesis of the guilt of Jones, and based on the pathological report and general knowledge about the world and causal relations within it, it follows deductively that Jones must have been near Smith within a specific time interval. If it can be established that Jones had an alibi for the time interval in question, then this new fact would be inconsistent with the assumption of his guilt. This is again based on general world knowledge that indicates that no object can be at two places at the same time. Ascertainment of the truth of Jones' alibi provides therefore more information than

any evidence that ties him to the dead, from his DNA on the crime scene to his fingerprints on the knife or the threats he uttered against Smith the night before. The investigating action that the system will recommend is the strength of Jones' alibi.

This work assumes that Jones either has or does not have an alibi. If he has a true alibi (*i.e.*, he was truly elsewhere at the time of the crime), pure logic dictates that the assumption of his guilt is falsified. It is then not necessary to calculate the relative weight of the evidence that supported the hypothesis that he was the perpetrator. Therefore, it is not necessary either to weigh the evidence that supports his alibi. The underlying world view is that of a *nave physicist*: the world consists of objects that impose causal effects on each other. It is possible to reason scientifically about these causal relations without any need of a mathematical framework. This approach has didactical, legal and computational advantages. Computationally, it keeps the system simple. Didactically, it transfers from the scientific sphere those modes of thought that prior analysis has found to be most desirable in the spheres of police work and court proceedings, without requiring more than that which is absolutely necessary for that purpose. Thinking in numbers does not come naturally to lawyers, and the mathematical nature of modern sciences seems to be one of the biggest stumbling blocks for interdisciplinary discourse [11]. Legally, assigning probabilistic values to individual pieces of evidence poses the danger of trespassing into the territory of the jury. By avoiding this type of evidence weighting altogether, lawyers could take the results provided by the system directly for the preparation of their case, without having to worry about procedural constraints.

However, the example of the alibi can also help to illustrate the self-imposed limitations of this approach, as the work is restricted in terms of practical relevance and applicability. Assume that there are again two competing explanations for the evidence found on the crime scene. According to the first case theory, Jones, a burglar, entered through a window overlooking the back garden, was surprised by the house owner Smith, stabbed him and killed him. According to the second theory, Miller and Smith entered into a fight over the rent. A struggle ensued during which Miller grabbed a kitchen knife and killed Smith. If the first theory is correct, there is a certain probability that pollen from the garden are on Jones' clothes. There is also a certain probability that fibres from his clothes are found on the window frame. If either pollen or fibres are found, the hypothesis is strengthened; their absence weakens the hypothesis. Similarly, if the second theory is correct, there is a chance that fibres from Miller can be found on Smith and fibres from Smith's clothes can be found on Miller. Consequently, there are four different tests that an investigating officer could request: pollen on Jones' clothes, fibres from Jones' clothes on the window frame, fibres from Miller's clothes on Smith, fibres from Smith's clothes on Miller. Which one should be prioritised? From which could the most be learned? Note that there are two different suspects, Jones and Miller, depending on which hypothesis is under investigation.

The system initially conceived could not answer these questions as they require reasoning about the respective strengths of different types of evidence. However, guidance to police officers in this sort of scenario has gained particular relevance through the introduction of actuarial methods in modern policing. As noticed above, police

investigations are subject to demands by various stake-holders that are often conflicting. Since the 1980s, managerial constraints on the police have increased considerably with the introduction of the internal market in forensic services [20, 21]. Cost constraints are therefore one of the parameters that can influence the decision making by an investigative officer. Yet, at the same time evidentiary constraints remain in place, and failure to carry out, for cost reasons alone, tests that have, scientifically, a good chance of exculpating a suspect may well jeopardize later prosecutions. It is therefore important that police officers are able to justify their investigating choices through sound scientific reasons, as much as managerial ones.

To enable the existing approach to support this type of decision making, it becomes necessary to increase its expressive power by combining the *assumption-based truth maintenance* technique [6] that has been used initially in building the software system with a representation capable of reasoning with probabilities. Assumption-based truth maintenance involves efficient tracking of the assumptions under which logical inferences are derived. This makes it easy to determine what may be inferred under which assumptions and what combinations of assumptions may lead to inconsistent reasoning results. The reminder of this paper shows in more technical detail how this is achieved.

It is important to remember though the context for which this work is developed. The intended application is to assist a police investigator to make, under uncertainty, decisions on how to reduce this uncertainty, how to optimise his so-called "information acquisition strategy". The primary aim is not to assist him in deciding whether or not the evidence against a suspect is sufficient, but in how to find that evidence in the first place. This is a categorically different scenario from that faced by evidence evaluation in court. See [13] for an example of network analysis for evidence evaluation and, also, the historical evidence available and not known, or used, at the time of the trial. There, all relevant investigating actions have typically been carried out, and the issue is to decide on the basis of the results of these actions if they prove the defendant's guilt to the required standards. In framing the problem considered herein as a question for information theory – how best to increase available information on the basis of already available evidence – basic concepts from information theory, in particular the concept of entropy, are adopted in developing the present work. As will be seen later, this approach to the investigation of a crime is not applicable for the evaluation of evidence because of its reliance on judgements about the probabilities of propositions.

The rest of this paper is organised as follows. The next section provides a brief overview of the underlying method taken for evidence investigation, setting the background for the present research. Section 3 shows the architecture of the proposed decision support system. Section 4 introduces the specific concepts employed and describes how different scenarios can be created systematically given evidence and generic domain knowledge. Section 5 presents the approach to analysing generated scenario spaces, with a focus on computing evidence collection strategies using entropy measures. The final section concludes the paper and points out further work. For convenience, the Appendix lists the notations used with explanations.

2 Bayesian Networks for Evidence Investigation

In order to produce effective evidence investigation strategies, a method is required to evaluate the informative value of a piece of evidence. As argued in [23], a wide range of methodologies have been devised for this purpose. The present work will employ Bayesian belief propagation to investigate 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, 24]. Therefore, this section presents a brief overview of this existing methodology.

Briefly, the method [8] for applying the Bayesian approach to evaluating a piece of forensic evidence e 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, a presumed "best defence", or a hypothesis of the forensic scientist or crime investigator.

3. Build a model to compute the probability $P(e \mid p_{\text{prosecution}})$ of the given piece of evidence being true 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*:

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

The ratio in (1) gives the value of the probability of the evidence if the prosecutions scenario is true relative to the probability of the evidence if the defence scenario is true. The vertical bar $\mid$ 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 defence) scenarios corresponds to a reasoning method that is widely adopted in science, and which is gaining increased acceptance in serious crime investigation [1]. The specific use of precise probabilities is more
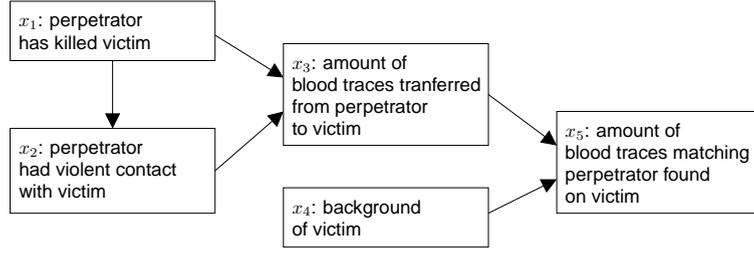
Figure 1: A simple Bayesian network

controversial, although it is adopted by major forensic laboratories, such as the Forensic Science Service of England and Wales [3]. 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 [16]. 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 [11].

The likelihood ratio method is crucially dependent upon a means to compute the probabilities $P(e \mid p_{\text{prosecution}})$ and $P(e \mid p_{\text{defence}})$. As shown in [2, 4], 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 causally dependent on H [18]. 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 $P(x_i \mid x_j, \ldots, x_k) : D_{x_i} \times D_{x_j} \times \ldots \times D_{x_k} \mapsto [0, 1]$ for each node x_i that has a set of (immediate) parent nodes $\{x_j, \ldots, x_k\}$ (where D_x, the domain of variable x, denotes the set of all values that can be assigned to x and $D_{x_i} \times \ldots \times D_{x_k}$ denotes the set of all combinations of values that can be assigned to $x_i, \ldots, x_k$).

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 unconditional 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. For example, 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 $P(x_2|x_1)$, $P(x_3|x_2, x_1)$ and $P(x_5|x_4, x_3)$ where the probabilities of the various possible combinations of outcomes are represented in tables. Table 1 is an example of such a probabilistic table:

$P(x_2|x_1)$	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: An example probabilistic table

Thus, given these figures, 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,

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

where, for example, $\sum_{v_2 \in D_{x_2}}$ denotes the sum over all possible values v_2 that x_2 may take. 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

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

$$P(x_2 : v_2 \mid x_1 : \text{false})$$

$$= \sum_{v_2 \in D_{x_2}} P(x_2 : v_2 \mid x_1 : \text{false}) \times \sum_{v_3 \in D_{x_3}} P(x_3 : v_3 \mid x_2 : v_2, x_1 : \text{false}) \times$$

$$\sum_{v_4 \in D_{x_4}} P(x_5 : \text{true} \mid x_3 : v_3, x_4 : v_4)$$

A BN can be employed as a means to calculate probabilities more efficiently. This increase in efficiency of the computation arises because the BN indicates which variables (characteristics which are measurements such as the elemental composition of glass) or factors (characteristics which are qualitative, such as whether or not the suspect had been in contact with the victim - true or false) 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
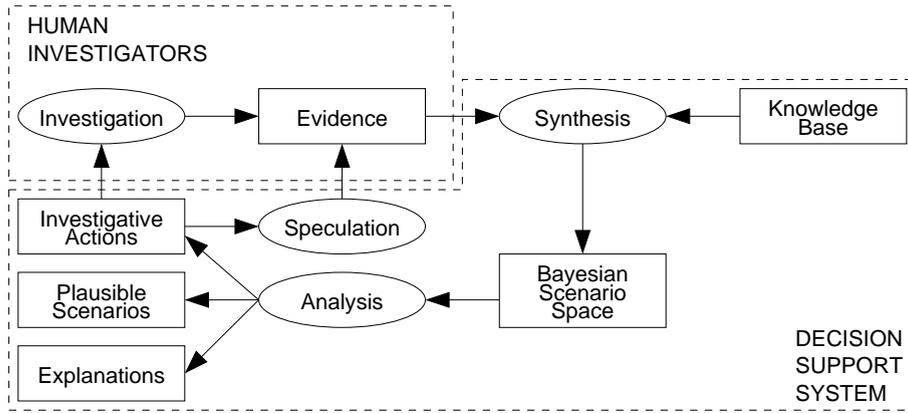
Figure 2: System Architecture

(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.

3 System Architecture

Evaluation of evidence for individual crimes requires the use of a specifically designed BN, and the use of BNs in the evaluation of evidence requires careful consideration for their construction. For crime investigation, in contrast to evidence evaluation, a method by which a generic BN may be constructed can be developed. The method is illustrated with an application to a situation in which a person is found dead from hanging. The method considers, first, the development of a novel inference mechanism that constructs a BN to describe a range of possible scenarios and corresponding hypotheses and, second, the employment of Bayesian model-based diagnostic techniques to evaluate evidence and to create strategies for the collection of evidence, with the payment of attention to the scenarios and hypotheses of the BN. The combination of the inference mechanism and the diagnostic techniques is known as a decision support system and the structure of the system is known as system architecture.

The overall architecture of the decision support system is shown in Figure 2. In this figure, types of information and knowledge are represented by rectangles, and inference mechanisms that reason with information/knowledge are represented by ellipses. 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 prespecified *knowledge base*. The scenario space describes how plausible states and events, pieces of evidence, investigating 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 inform-

ation, 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) investigating actions that are likely to yield useful information. The proposed investigating actions can, in turn, be used to inform the investigation directly or to *speculate* about the possible outcomes of the investigating action, determine future actions, and, hence, gather further evidence for input to a new synthesis operation.

4 Generation of a Space of Plausible Scenarios

4.1 Stored Knowledge

4.1.1 Situations

Consider the suspicious circumstance of the discovery of a dead body hanging by the neck. There are various scenarios which could have given rise to this circumstance and these are discussed in detail later. Whilst the investigation is continuing, these scenarios are grouped together in what is known as a *space* of plausible scenarios. The discussion of such scenarios involves what are known as *situations*. A situation is either a particular condition or status of the world, or a development that changes another situation.

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 (true or false with reference to a statement such as 'X killed Y') are also an important feature of situations. In a manner similar to the distinction between types and instances, the system may sometimes be interested in situations that denote specific quantities and sometimes it may not. Situations 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*, where the quantity of fibres transferred is left unspecified. 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 one of three values 'none', 'few' or 'many', for example (though other classifications are possible).

Each situation variable x is said to have a *domain D_x*. This domain is the set of quantities that the situation variable can take in a valid assignment. All situations of the same type are given the same domain, which is defined in the knowledge base. The assignment of a domain value to a variable may be affected by multiple influences. The way in which the effects of different influences on the assignment of a variable x are to be combined is defined by a *combination operator $\oplus_x$*, which is in general specific to that variable. As with the domains, combination operators are associated with variable types in the knowledge base.

Note that in what follows, the terms *situations* may interchangeably refer to variable types, variable instances, assignment types or assignment instances. Whenever the context does not clarify any ambiguity between these concepts, however, the text will specify to which context reference is being made.

Most situations used simply describe real-world situations. However, some also convey additional information that may aid in decision support. These concepts have been adapted from earlier work on abductive reasoning [19] and model based diagnosis [12]. In particular, certain situations may 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 herein defined as "information", it does not equal the "exhibits" presented in court. Thus, for example, in the context of this discussion, the existence of a suicide note is not considered to be a piece of evidence on its own, but the conclusions of a handwriting expert who has analysed the note are. Different from evidence, *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 on possible justifications. For instance, when a person is given the responsibility of analysing the cause of death of a victim, the status of that person as a medical expert is normally deemed to be a fact (subject to current issues concerning the status of expert witnesses which are not part of the thesis of this paper). *Hypotheses* are possible answers to questions that are 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. Three types of assumption are considered here: (i) *Investigating actions* which 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 investigating action. Each investigating action is associated with an exhaustive set of mutually exclusive outcomes for the consideration of a piece of evidence that cover all possible situations (exhaustive outcomes) considered by the action of the investigation, such that there can be one and only one outcome of the investigation of the piece of evidence (*e.g.*, the blood group of a suspect). (ii) *Default assumptions* which are presumed true unless they are contradicted. Such assumptions are typically employed to represent the conditions from which an expert produces evaluations based upon sound methodology and understanding of his/her field. (iii) *Conjectures* which simply correspond to uncertain states and events that need not be described as consequences of other states and events.

4.1.2 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 so-called *domain knowledge* that is relevant to a crime investigation concerns the causal relations among the plausible situations.

This is 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.

The rule component of a scenario fragment is an expression of the form:

IF   *conjunction of antecedents*

ASSUMING   *conjunction of assumptions*

THEN   *consequent*

where antecedents, assumptions and consequent refer to 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. Note that the assumptions are situations that may be presumed to be present in a scenario because they refer to conjectures, default assumptions or investigating 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 following rule R describes the phenomena that persons who are addicted to a certain substance have traces of that substance in their blood, especially when they have recently used that substance.

IF   anaesthetic A is an addictive substance (a_1)

ASSUMING   person P has an addiction to anaesthetics (a_2), and

person P has recently used anaesthetic A (a_3)

THEN   person P has anaesthetic A in blood (c)

A scenario fragment (SF) consists of a rule and a set of probability distributions. For example, the rule could have variable types $a_1, \ldots, a_m$ as its antecedents and assumptions and the variable type c as its consequent. The SF then includes probability distributions over the possible assignments of c, for different combinations of assignments to the variables $a_1, \ldots, a_m$. These distributions describe how the assignment of c is affected by the phenomena under different configurations of $a_1, \ldots, a_m$. In what follows, the probability distributions of such a scenario fragment SF will be denoted by:

$$P(a_1 : v_1, \ldots, a_m : v_m \xrightarrow{SF} c : v_c)$$

where $a : v$ denotes the situation assignment obtained by assigning v to situation variable a, and $a_1 : v_1, \ldots, a_m : v_m, c : v_c$ are assignments corresponding to the variables $a_1, \ldots, a_m, c$. This can be applied to the above example rule R. Let SF denote the scenario fragment containing this rule with situation variable types a_1, a_2, a_3, c. Let a_1, a_2, a_3 be variables that are assigned one value from {true, false} (indicating whether the condition is met or not) and c be a variable that is assigned one of {high, low, none} (indicating the proportion of the addictive substance in the person's blood). Then, the following probability distributions associated with scenario fragment SF may be:

$$P(a_1 : \text{true}, a_2 : \text{true}, a_3 : \text{true} \xrightarrow{SF} c : \text{high}) = 0.8$$
$$P(a_1 : \text{true}, a_2 : \text{true}, a_3 : \text{true} \xrightarrow{SF} c : \text{low}) = 0.2$$
$$P(a_1 : \text{true}, a_2 : \text{true}, a_3 : \text{true} \xrightarrow{SF} c : \text{none}) = 0$$

with

$$P(a_1 : v_1, a_2 : v_2, a_3 : v_3 \xrightarrow{SF} c : \text{high}) = 0$$
$$P(a_1 : v_1, a_2 : v_2, a_3 : v_3 \xrightarrow{SF} c : \text{low}) = 0$$
$$P(a_1 : v_1, a_2 : v_2, a_3 : v_3 \xrightarrow{SF} c : \text{none}) = 1$$

for any combination of values $(v_1, v_2, v_3) \neq (\text{true}, \text{true}, \text{true})$. These distributions indicate that, if the antecedent and assumption conditions are met, the phenomenon will cause high levels of anaesthetic A in the blood of person P with probability 0.8 and low levels with probability 0.2, and, naturally, no anaesthetic with probability 0, and that the phenomenon is not in effect if the conditions are not met (though this does not exclude the possibility that other phenomena may still cause high levels of anaesthetic A in the blood of person P).

4.1.3 Presumptions

The use of scenario fragments in the construction of BNs is enabled by assuming that SFs in the knowledge base possess the following properties:

1. *Any two probability distributions from two scenario fragments that involve 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 consequent variable 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 variable instance to be affected by itself; a consequent variable cannot be its own antecedent of assumption. This is required because BNs cannot represent such information as they are inherently acyclic [17].

While presumption 1 is a strong assumption, and may hence reflect a significant limitation of the present work, it is required to compute efficiently the combined effect of a number of scenario fragments on a single variable (see Section 4.4). Future work will seek to relax this assumption in order to generalise further the application of the method proposed.

4.2   Scenario Spaces

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 (which allows efficient tracking of the assumptions under which logical inferences are derived) [6], the structural scenario space maintains the smallest sets of other situations that can meaningfully explain it.

  Consider, for example, a crime scene containing the dead body of a man, who is named johndoe in what follows, hanging from a rope. Assume that near the body of johndoe, a knife is found. Different scenarios may provide a valid explanation for the dead body and the discovery of the knife. For instance:

  - johndoe may have hung himself as part of an autoerotic asphyxiation ritual with the intention of cutting himself loose with the knife, but failed to do so; or

  - johndoe may have been hung by a person who overpowered him using the knife.
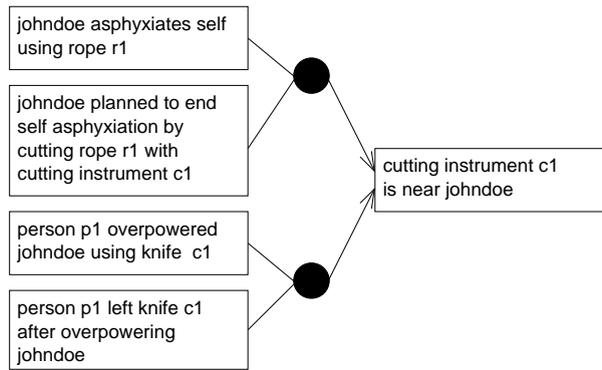
Figure 3: Partial scenario space

Figure 3 is a partial structural scenario space that contains some of the information relevant to these possible scenarios. In particular, the 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".

In Figure 3, the black circles correspond to the keyword *and* in the above description of causes. Thus, the black circles combine multiple causes that, in conjunction, are sufficient to explain the consequent cutting instrument c1 is near johndoe.

Note that in all following figures, black circles represent the same logical conjunctive operator *and* as above. The cases where the justification only involves one cause node (e.g., the link from johndoe committed suicide by hanging to johndoe died from hanging) are simply a degeneralised form of this general representation (in the sense of the number of causes being one). In these cases, a black circle remains in the middle of the justification link, as these figures are produced automatically.

- *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 of which Table 1 is an example.

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 minim-

ally explanatory scenarios, it does, as shown in Section 2, provide the necessary knowledge to test hypotheses and evaluate evidence. The remainder of this section aims to demonstrate how both scenario spaces can be constructed from a given knowledge base. This demonstration will start from the partial scenario space in Figure 3.

4.3 Generation of Structural Scenario Spaces

The approach to generate the structure of a scenario space relies on the application of two conventional inference techniques to the knowledge base. They are (i) abduction of the plausible causes of known or hypothesised situations, and (ii) deduction of the plausible consequences of known or hypothesised situations. This is basically the same as reported in the initial system [22]. However, for completeness, a brief summary of the relevant techniques is given below, with illustrations.

The abduction and deduction operations are applied to develop the scenario space by adapting the rules from the knowledge base to given situations incrementally. A partially constructed scenario space is termed an emerging scenario space during such an abduction and deduction process:

- *Abduction*: given a piece of information that matches the consequent of a rule 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. For example, given the rule:

    IF   anaesthetic A is an addictive substance

    ASSUMING   person P has an addiction to anaesthetics, and

        person P has recently used anaesthetic A

    THEN   person P has a high level of anaesthetic A in blood

    and a scenario space that contains the piece of information $n_1 = $ "person vic has a high level of anaesthetic percocet in blood", then the abduction operation will generate the pieces of information:

    n_2 = "anaesthetic percocet is an addictive substance"

    a_1 = "person vic has an addiction to anaesthetics"

    a_2 = "person vic has recently used anaesthetic percocet"

    and add the implication $n_2 \wedge a_1 \wedge a_2 \rightarrow n_1$ to the emerging scenario space. Notationally, throughout this work, those nodes that represent assumptions are denoted by a_i (e.g., the a_1 and a_2 above), and those nodes that represent either facts or consequents derived from other nodes are denoted by n_j (e.g. the n_1 and n_2 above).

- *Deduction*: given a set of pieces of information that match the antecedents of a rule 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. For example, given the rule:

15

IF	person I is an investigator, and	
	person P has symptom S, and	
	symptom S corresponds to purple spots on the eyes, and	
	symptom S is petechiae on the eyes	
ASSUMING	investigator I correctly diagnoses symptom S	
THEN	investigator I identifies symptom S as petechiae on the eyes of person P	

and a scenario space that contains the pieces of information $n_1 = $ "person quincy is an investigator", $n_2 = $ "person vic has symptom x", $n_3 = $ "symptom x corresponds to purple spots on the eyes" and $n_4 = $ "symptom x is petechiae on the eyes", then the deduction operation will generate the pieces of information:

a_1 = "investigator quincy correctly diagnoses symptom x"

n_5 = "investigator quincy identifies symptom x as petechiae on the eyes of person vic"

and add the implication $n_1 \wedge n_2 \wedge n_3 \wedge n_4 \wedge a_1 \rightarrow n_5$ to the emerging scenario space.

By means of such abduction and deduction steps, a structural scenario space can be created from a given set of available evidence, using the following procedure:

1. *Initialisation*: A new structural scenario space is first created containing only the pieces of evidence and no inferences between them. For simplicity, this procedure will be illustrated using an example involving just a single piece of evidence that represents the "hanging body of johndoe".

2. *Abduction of plausible causes from collected evidence*: Abduction involves speculation about the possible causes of a known situation. Abduction operations are applied with all the scenario fragments in the knowledge base, whose consequent matches a situation variable instance in the emerging scenario space. Obviously, as these abduction operations are applied (initially for those scenario fragments whose consequent matches a piece of evidence), new situation variable instances are added to the emerging scenario space, and therefore, new scenario fragments become applicable. In so doing, causal chains of situations are created, and these constitute plausible explanations for the available evidence. This part of the procedure ends when no more information can be adduced and added to the scenario space.

Figure 4 depicts a sample emerging scenario space that may be the outcome of this step (depending on exactly what is stored in the knowledge base of course). For example, given a rule

IF	P died from hanging
THEN	hanging body of P

and the initial piece of evidence "hanging body of johndoe", the abduction step adds the node "johndoe died from hanging" as well as the implication:
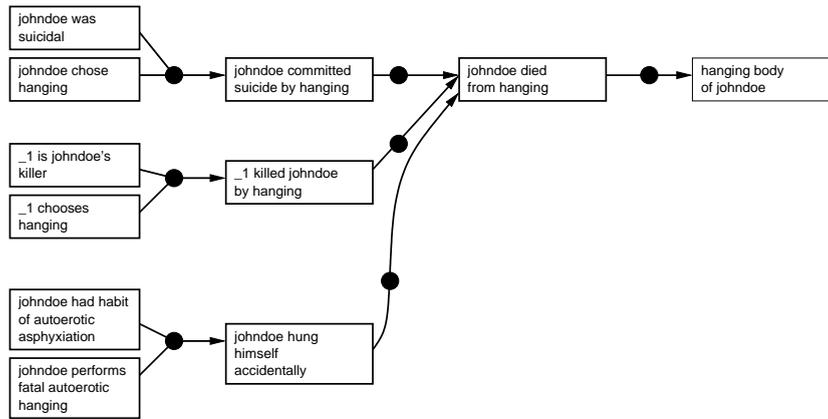
Figure 4: Sample scenario space after abduction of plausible causes from collected evidence

$$\text{johndoe died from hanging} \rightarrow \text{hanging body of johndoe}$$

The other nodes of Figure 4 are created in a similar manner (by backward application of the rules in the knowledge base).

3. *Deduction of plausible consequences from plausible causes*: Deduction operations are applied with all the scenario fragments in the knowledge base, whose antecedents match a set of situation variable instances in the emerging scenario space. As with the abduction process, new information that is deduced is added to the emerging scenario space and may give rise to new matches of scenario fragment antecedents with information in it. 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.

Figure 5 depicts a sample emerging scenario space that may be the outcome of this step. In this figure, the gray squares and bold arrows respectively denote the nodes and justifications that have been added by the procedure that expanded the emerging scenario space of Figure 4 to the one shown here. A similar treatment is given to the next figure.

4. *Abduction of plausible causes from uncollected evidence*: When assessing the value of searching for additional evidence through further investigating 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
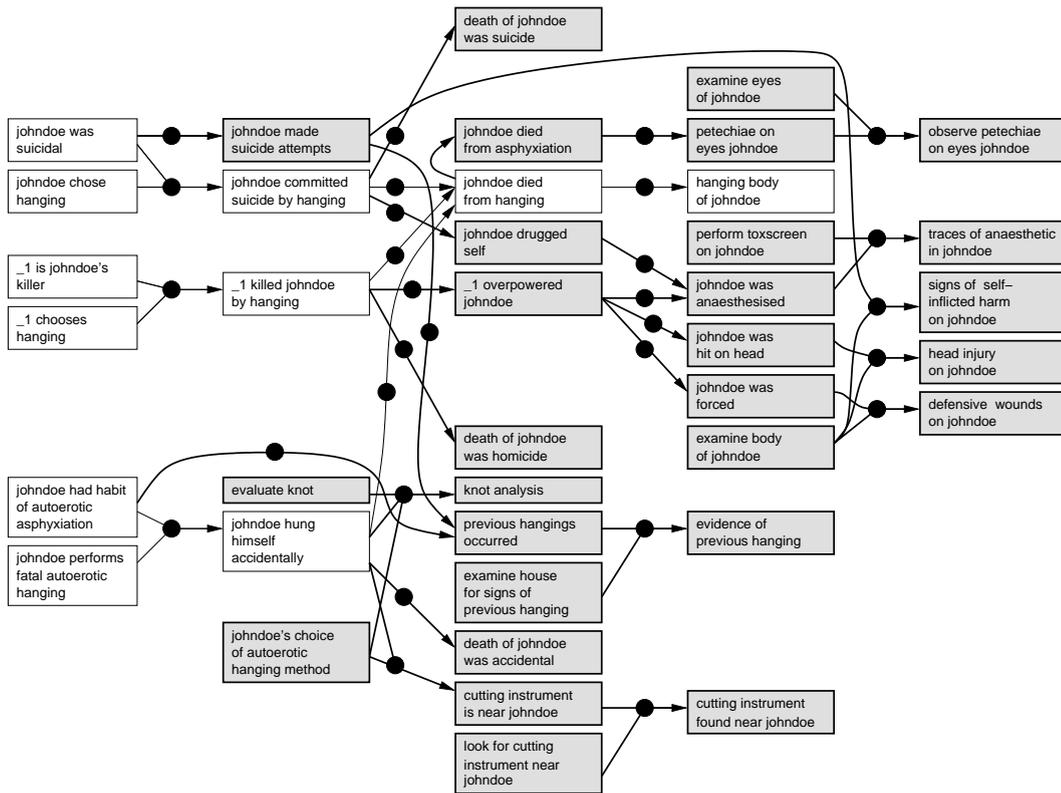
Figure 5: Sample scenario space after deduction of plausible consequences from plausible causes

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. Figure 6 depicts a sample emerging scenario space that may be the outcome of this step.

5. *Deduction of inconsistencies*: Once a structural 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 structural scenario space. This involves applying a deduction operation for each inconsistency whose situation variables match a set of situation variables in the emerging structural scenario space. Figure 7 depicts a sample emerging scenario space that may be the outcome of this step, where each pair of nodes whose links are merged into the specific node nogood represents an inconsistency (in other words, nogood denotes the collection of any given and deduced inconsistencies). Note that in this figure, for easy reference, each node is labelled with a node number; e.g., the node johndoe was suicidal is labelled with n_1 and defensive wounds on johndoe with n_{39}.

For example, given the inconsistency
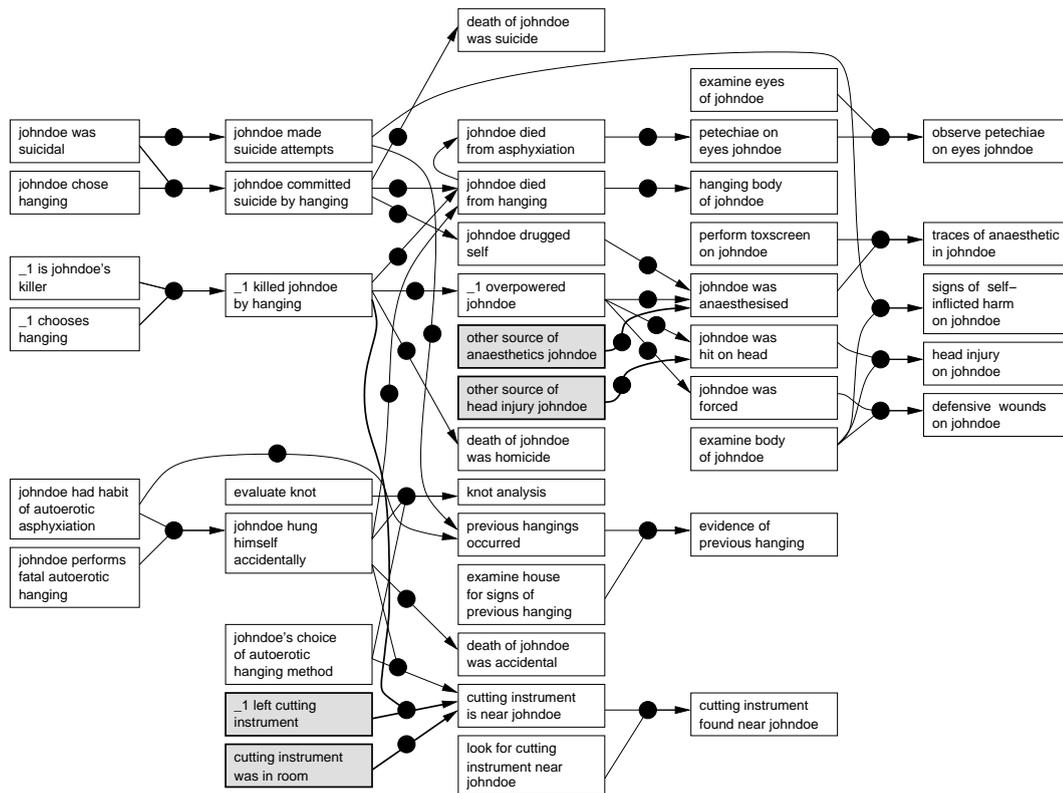
Figure 6: Sample scenario space after abduction of plausible causes from uncollected evidence

P committed suicide by hanging : TRUE, and

P hung himself accidentally : TRUE

and assuming that the emerging scenario space contains the hypothetical pieces of information "johndoe committed suicide by hanging" and "johndoe hung himself accidentally", the deduction of inconsistencies step adds the implication

johndoe committed suicide by hanging $\wedge$ johndoe hung himself accidentally $\rightarrow$ nogood

to the emerging scenario space.

The structural scenario space generated by the above procedure enables the decision support system to trace back the conjectures and facts that ultimately could 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 situations and that are not assumptions or facts. These are spurious explanations, and they are to be removed from the emerging scenario space. Therefore, the procedure is extended with a sixth step:

6. *Spurious explanation removal*: Any situation variable instance in the emerging scenario space that is not the consequent of one or more other situation instances and that is neither an assumption nor a fact is removed.

Figure 7: Sample scenario space after deduction of inconsistencies

Consider, for instance, the following rule:

> IF   person W is in location L1, and
>
> person V is assaulted in location L2, and
>
> location L1 is near location L2
>
> ASSUMING   person V screams due to the assault
>
> THEN   person W witnesses a scream from V

Given this rule and an observation: person rita witnesses a scream from johndoe, the abduction operation will generate the following plausible explanation:
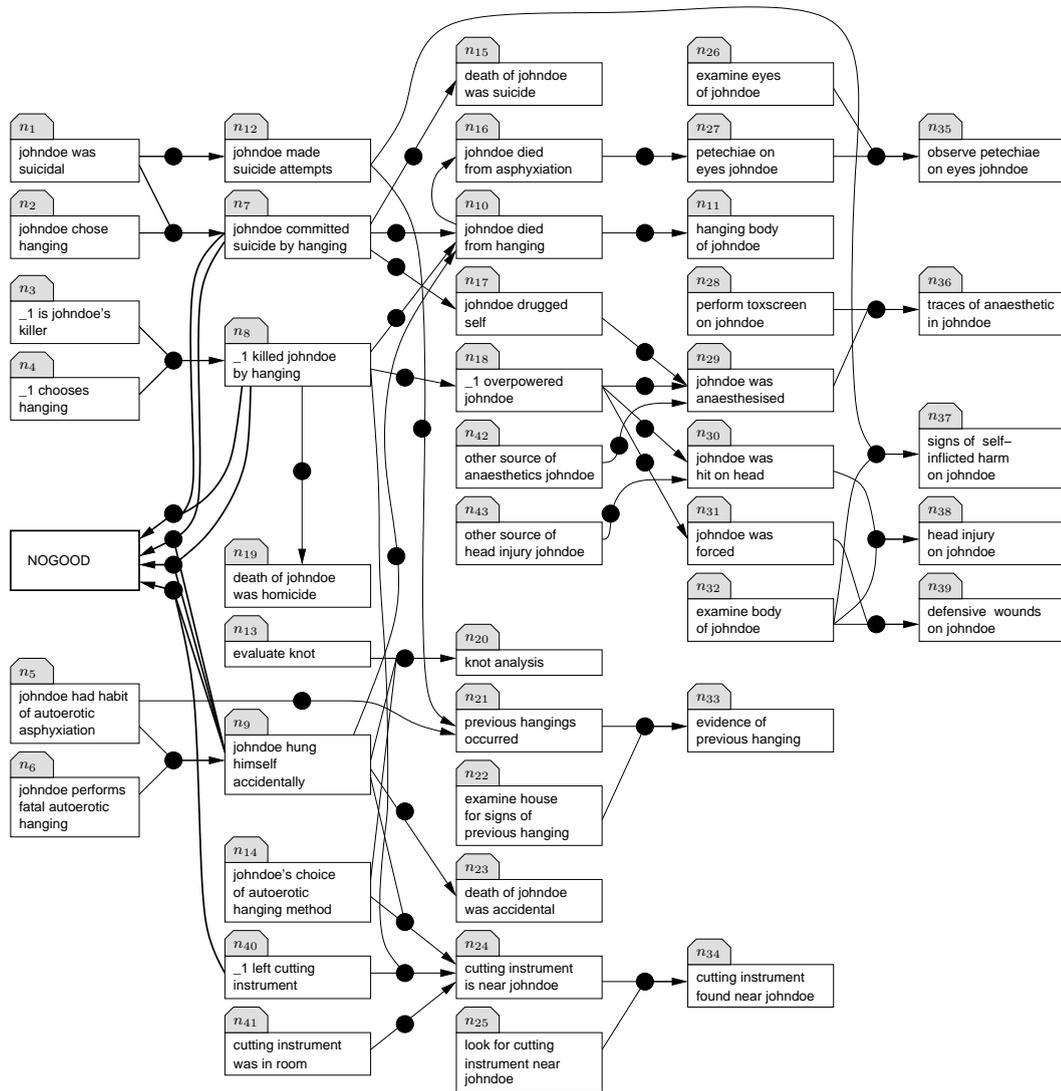
> person rita is in location lounge, and
>
> person johndoe is assaulted in location garden, and
>
> location lounge is near location garden
>
> person johndoe screams due to the assault

Note, however, that location lounge is near location garden cannot be guaranteed to be true (here, for illustrative simplicity, presuming that person rita is in location lounge and person johndoe is assaulted in location garden are either given observations or have been inferred by a previous deduction/abduction process). If this is the case (and assuming that the variable is modelled to have a boolean or {true, false} domain), this piece of information may not be asserted by the investigator. As such, the piece of information location lounge is near location garden is neither a fact nor an assumption and this plausible explanation is therefore removed from the emerging scenario space.

4.4   Generation of Probabilistic Scenario Spaces

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 $\ldots \wedge a \wedge \ldots \rightarrow c$, 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) are 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 last step of this procedure may appear not to be straightforward, therefore requiring further explanation. Thus, the remainder of this section describes how such conditional probability distributions are calculated.

Let $\{p_1, \ldots, p_s\}$ 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 $\{p_1 : v_1, \ldots, p_s : v_s\}$, 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 $p_1 : v_1, \ldots, p_s : v_s$ affect p_c. Let $SF_1, \ldots, SF_k$ be such scenario fragments and $p_c : c_1, \ldots, p_c : c_k$ be their respective outcomes. p_c is assigned c whenever the combined effect of the scenario fragments $c_1 \oplus \ldots \oplus c_k = c$, with regard to a certain predefined interpretation of the combination operator $\oplus$. Many operators may be adopted to implement this [25]. In a given application, the choice of what combination operator to use requires careful consideration of several important factors such as computational efficiency, intuitive interpretability and the power of such an operator in aggregating information.

From the above, the conditional probability $P(p_c : c \mid A)$ is given by:

$$P(p_n : c \mid A) = P\left[\bigvee_{c_1 \oplus \ldots \oplus c_k = c}\left(\bigwedge_{i=1,\ldots,k}\left(A \overset{SF_i}{\to} p_n : c_i\right)\right)\right] \tag{4}$$

where $\bigvee$ and $\bigwedge$ represent logical disjunction and conjunction (i.e., "or" and "and"), respectively.

Thus, computing $P(p_n : c \mid A)$ involves calculating the likelihood of a combination of events described by a so-called disjunctive normal form expression (i.e. a logical expression of the form $(x_{11}$ and $\ldots$ and $x_{1m_1})$ or $\ldots$ or $(x_{n1}$ and $\ldots$ and $x_{nm_n}))$. Because the occurrence of the different combinations of outcomes $c_1, \ldots, c_k$ of scenario fragments $SF_1, \ldots, SF_k$ involves mutually exclusive events, the calculation can be resolved by adding the probabilities of the conjuncts in (4):

$$P(p_n : c \mid A) = \sum_{c_1 \oplus \ldots \oplus c_k = c} P\left(\bigwedge_{i=1,\ldots,k}\left(A \overset{SF_i}{\to} p_n : c_i\right)\right) \tag{5}$$

From presumption 1 (as noted in 4.1.3), 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 (5) is equal to the product of the probabilities of their conjuncts, thereby

$$P(p_n : c \mid A) = \sum_{c_1 \oplus \ldots \oplus c_k = c} \left(\prod_{i=1,\ldots,k} P\left(A \overset{SF_i}{\to} p_n : c_i\right)\right) \tag{6}$$

Consider the following two scenario fragments, which form part of the probabilistic scenario space that corresponds to the structural space of Figure 7:

     IF   person V had an auto-erotic asphyxiation habit (n_5)

   THEN   person V has previously hung him/herself (n_{21})

     IF   person V has attempted to commit suicide (n_{12})

   THEN   person V has previously hung him/herself (n_{21})

where n_5 and n_{12} have a boolean domain $\{\top, \bot\}$ (respectively denoting true, false), and n_{21} to a variable taking values from a (totally ordered) domain, say $\{\texttt{never}, \texttt{veryfew}, \texttt{several}\}$, with $\texttt{never}<\texttt{veryfew}<\texttt{several}$. Here, purely for illustration, these three values are used. For different applications a different number of different terms may be employed. Also for illustration purposes, the combination operator associated with the last variable is set to $\max$ (which enjoys computational efficiency). Then, the probabilities of assignments to n_{21} given $n_5 : \top$ and $n_{12} : \top$, can be computed as follows. For notational convenience, let the above two scenario fragments be named SF_1 and SF_2. Assume that the unconditional prior probabilities associated with scenario fragments are:

$$P(n_5 : \top \overset{SF_1}{\to} n_{21} : \texttt{never}) = 0.10$$

$$P(n_5 : \top \overset{SF_1}{\to} n_{21} : \texttt{veryfew}) = 0.40$$

$$P(n_5 : \top \overset{SF_1}{\to} n_{21} : \texttt{several}) = 0.50$$

$$P(n_{12} : \top \overset{SF_2}{\to} n_{21} : \texttt{never}) = 0.70$$

$$P(n_{12} : \top \overset{SF_2}{\to} n_{21} : \texttt{veryfew}) = 0.29$$

$$P(n_{12} : \top \overset{SF_2}{\to} n_{21} : \texttt{several}) = 0.01$$

where, for instance, $P(n_5 : \top \overset{SF_1}{\to} n_{21} : \texttt{never}) = 0.10$ represents the statement that the probability of which scenario fragment SF_1 assigns the value $\texttt{never}$ to variable n_{21} given n_5 being true is 0.10.

Following (6), to compute $P(n_{21} : \texttt{veryfew}|n_5 : \top, n_{12} : \top)$ all combinations of the outcomes c_1 and c_2 of scenario fragments SF_1 and SF_2, with $c_1, c_2 \in \{\texttt{never}, \texttt{veryfew}, \texttt{several}\}$, such that $\max(c_1, c_2) = \texttt{veryfew}$, must be considered. This is because (6) states that the probability of n_{21} taking on the value of $\texttt{veryfew}$, given that n_5 and n_{12} are both true, is the sum of the products of the probabilities of any two possible SF assignment outcomes, subject to the constraint that the composition (or maximum in this case) of both outcomes is $\texttt{veryfew}$. There are three such combinations: $\{c_1 : \texttt{veryfew}, c_2 : \texttt{veryfew}\}$, $\{c_1 : \texttt{never}, c_2 : \texttt{veryfew}\}$ and $\{c_1 : \texttt{veryfew}, c_2 : \texttt{never}\}$. Hence,

$$P(n_{21} : \texttt{veryfew}|n_5 : \top, n_{12} : \top)$$
$$= P(n_5 : \top \overset{SF_1}{\to} n_{21} : \texttt{veryfew}) \times P(n_{12} : \top \overset{SF_2}{\to} n_{21} : \texttt{veryfew}) +$$
$$P(n_5 : \top \overset{SF_1}{\to} n_{21} : \texttt{never}) \times P(n_{12} : \top \overset{SF_2}{\to} n_{21} : \texttt{veryfew}) +$$
$$P(n_5 : \top \overset{SF_1}{\to} n_{21} : \texttt{veryfew}) \times P(n_{12} : \top \overset{SF_2}{\to} n_{21} : \texttt{never})$$
$$= 0.4 \times 0.29 + 0.1 \times 0.29 + 0.4 \times 0.7 = 0.425$$

n1 — johndoe was suicidal

v_1	true	false
$P(n_1 : v_1)$	0.02	0.98

n2 — johndoe chose hanging

v_2	true	false
$P(n_2 : v_2)$	0.1	0.9

n3 — _1 is johndoe's killer

v_3	true	false
$P(n_3 : v_3)$	0.01	0.99

n4 — _1 chooses hanging

v_4	true	false
$P(n_4 : v_4)$	0.05	0.95

n5 — johndoe had habit of autoerotic asphyxiation

v_5	true	false
$P(n_5 : v_5)$	0.01	0.99

n6 — johndoe performs fatal autoerotic hanging

v_6	true	false
$P(n_6 : v_6)$	0.025	0.975

n7 — johndoe committed suicide by hanging

$P(n_7 : v_7 | n_1 : v_1, n_2 : v_2)$

		v_7	
v_1	v_2	true	false
true	true	1	0
true	false	0	1
false	true	0	1
false	false	0	1

n8 — _1 killed johndoe by hanging

$P(n_8 : v_8 | n_3 : v_3, n_4 : v_4)$

		v_8	
v_3	v_4	true	false
true	true	1	0
true	false	0	1
false	true	0	1
false	false	0	1

n9 — johndoe hung himself accidentally

$P(n_9 : v_9 | n_5 : v_5, n_6 : v_6)$

		v_9	
v_5	v_6	true	false
true	true	1	0
true	false	0	1
false	true	0.05	0.95
false	false	0	1

n10 — johndoe died from hanging

$P(n_{10} : v_{10} | n_7 : v_7, n_8 : v_8, n_9 : v_9)$

			v_{10}	
v_7	v_8	v_9	true	false
true	true	true	1	0
true	true	false	1	0
true	false	true	1	0
true	false	false	1	0
false	true	true	1	0
false	true	false	1	0
false	false	true	1	0
false	false	false	0	1

n11 — hanging body of johndoe

$P(n_{11} : v_{11} | n_{10} : v_{10})$

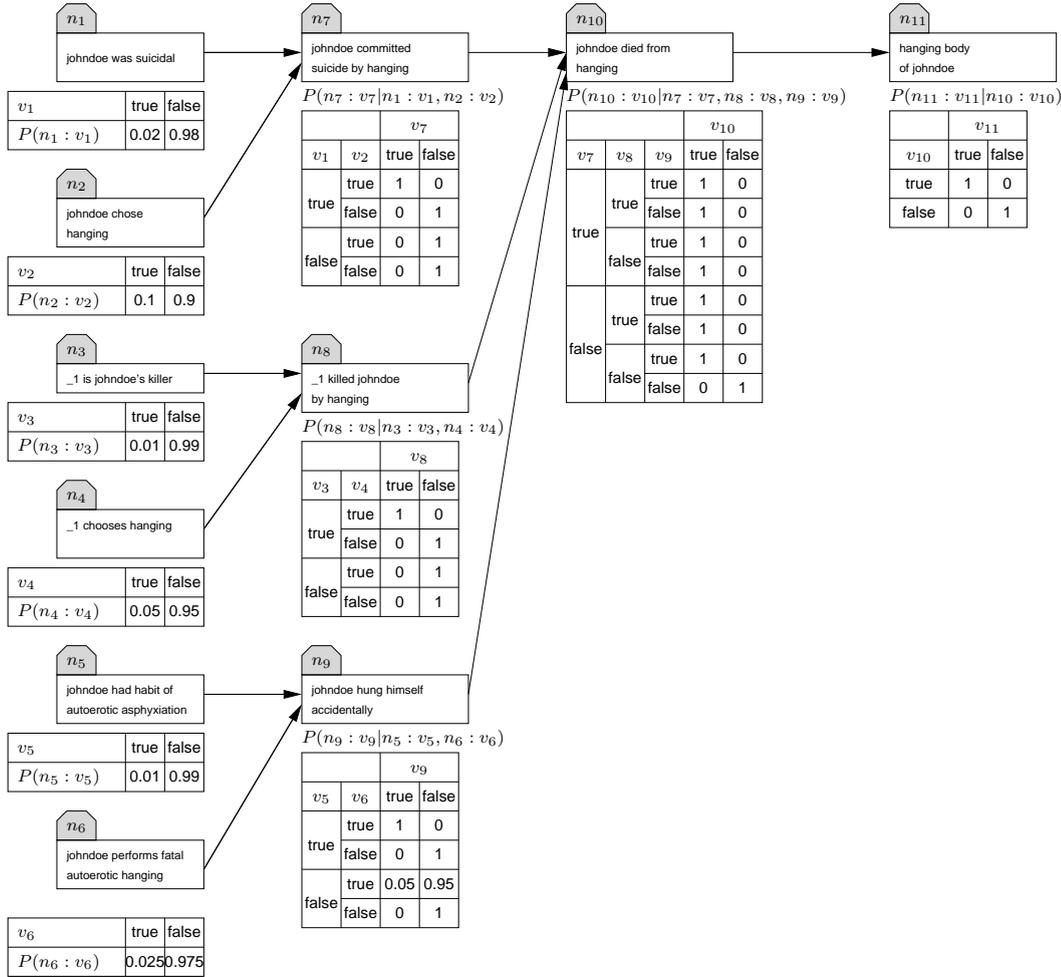| v_{10}	v_{11} true	false
true	1	0
false	0	1

Figure 8: Sample partial Bayesian network

Similarly, it can be shown that

$$P(n_{21} : \texttt{never} | n_5 : \top, n_{12} : \top) = 0.070$$

$$P(n_{21} : \texttt{several} | n_5 : \top, n_{12} : \top) = 0.505$$

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 8 for illustration, which relates to the corresponding part of the scenario space of Figure 7.

5 Analysis of a Generated Space of Plausible Scenarios

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 particular variable assignments. For implementational simplicity, hypotheses and investigating actions are assumed to be represented by (truth) assignments to boolean variables (although this can be extended).

5.1 Hypothesis sets and query types

The approach shown herein 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:

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

Note that in implementing this work, hypothesis sets are given 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 the victim in a suspicious death case, or the killer of a victim in a homicide case.

5.2 Entropy

As indicated previously, the work here employs an information theory based approach, which is widely used in areas such as machine learning [15] and model based diagnosis [12]. Information theory utilises the concept of *entropy*, which is a measurement of how broadly doubt is distributed over a range of choices. Applied to the present problem, the entropy over an exhaustive set of m 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 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 $log(m)$ (which corresponds to a total lack of information) when $P(h_1) = \ldots = P(h_m) = \frac{1}{m}$, and that $\epsilon(H)$ reaches 0, its lowest value (which corresponds to a totally certain situation) when all $P(h_i) = 0$, $i = 1, \ldots, j-1, j+1, \ldots, m$ and $P(h_j) = 1, j \in \{1, \ldots, m\}$.

In crime investigation, additional information is created through evidence collection. The entropy, given evidence E, known as the entropy metric, is used for the purpose of generating evidence collection strategies. The entropy metric 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:

$$\epsilon(H \mid E) = -\sum_{h \in H} P(h \mid E) \log P(h \mid E) \tag{7}$$

where the values $P(h \mid E)$ (the probability of the truth of the hypothesis given the evidence) can be computed through BN inference. This expression may be considered as a measure of the lack of information about H given E. For the example problem from the sample scenario space, the following probabilities can be computed, with E_1 containing hanging body of johndoe : $\top$ (which corresponds to the available evidence) and nogood : $\bot$ (which represents the absence of inconsistencies):

$$P(\text{death of johndoe was homicide} \mid E_1) = 0.22$$
$$P(\text{death of johndoe was suicide} \mid E_1) = 0.33$$
$$P(\text{death of johndoe was accidental} \mid E_1) = 0.45$$
$$P(\text{death of johndoe was natural} \mid E_1) = 0 \tag{8}$$

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 investigating actions a from a given set A according to the criterion of minimising the expected lack of information about H given E:

$$\min_{a \in A} Exp(\epsilon(H \mid E), a) \tag{9}$$

where $Exp(x)$ denotes the expected value (or expectation) of random variable x, which is defined to be the sum of the probability of each possible outcome of x multiplied by its payoff (that is normally termed its "value"). In other words, $Exp(x)$ represents the average amount that x is expected to be if x is repeatedly measured many times.

Note that the entropy values calculated by (7) are affected by the prior distributions assigned to the assumptions, as described in Section 4.1. 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 instance, it may be the case that the probability of suicides or autoerotic deaths may be underestimated by inexperienced investigators. As such, the strategy dictated by (9) offers a useful means to decide on what evidence to collect next. Nevertheless, information yielded by the minimal entropy decision rule should not be used for evidence evaluation in court.

5.3 Minimal entropy-based evidence collection

Let a denote an investigating 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 investigating 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):

$$Exp(\epsilon(H \mid E), a) = \sum_{e \in E_a} P(e \mid a : \top, E)\epsilon(H \mid E \cup \{a : \top, e\}) \tag{10}$$

where $\top$ refers to true. The ongoing example contains an investigating action $a = $ perform toxscreen on johndoe: $\top$, representing a toxicology test of johndoe searching for traces of anaesthetics and a corresponding set of outcomes $E_a = \{$traces of anaesthetic in johndoe $: \top,$ traces of anaesthetic in johndoe $: \bot\}$, respectively denoting a positive toxscreen and a negative one, where $\bot$ refers to false. Let E_2 be a set containing hanging body of johndoe: $\top$, and nogood: $\bot$. Then, through exploitation of the probabilistic scenario space the following can be computed:

$$P(\text{traces of anaesthetic in johndoe} : \top \mid E_2) = 0.17$$

$$P(\text{traces of anaesthetic in johndoe} : \bot \mid E_2) = 0.83$$

$$P(\text{death of johndoe was homicide} \mid E_2 \cup \{\text{traces of anaesthetic in johndoe} : \top\}) = 0.40$$

$$P(\text{death of johndoe was suicide} \mid E_2 \cup \{\text{traces of anaesthetic in johndoe} : \top\}) = 0.49$$

$$P(\text{death of johndoe was accidental} \mid E_2 \cup \{\text{traces of anaesthetic in johndoe} : \top\}) = 0.11$$

$$P(\text{death of johndoe was homicide} \mid E_2 \cup \{\text{traces of anaesthetic in johndoe} : \bot\}) = 0.19$$

$$P(\text{death of johndoe was suicide} \mid E_2 \cup \{\text{traces of anaesthetic in johndoe} : \bot\}) = 0.43$$

$$P(\text{death of johndoe was accidental} \mid E_2 \cup \{\text{traces of anaesthetic in johndoe} : \bot\}) = 0.38$$

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 for traces of anaesthetic can be computed as the following instance:

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

The investigating 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 investigating 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 investigating actions in this case are therefore established to be a search for a cutting instrument and a knot analysis. This result can be intuitively explained by the fact that these investigating 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.

5.4 Extensions

The approach presented above offers a possible alternative to the conventional likelihood ratio approach. Several further improvements of this technique are proposed here. As these extensions are not difficult to conceive, detailed

examples are omitted below.

5.4.1 Local optima and action sequences

Although the minimum EPE evidence collection technique guarantees to return an effective investigating 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 $a_1, \ldots, a_v$ (of course, with incurred overheads over computation and potential difficulty in considering multiple actions at one time):

$$
\begin{aligned}
&Exp(\epsilon(H \mid E), a_1, \ldots, a_v) \\
&= \sum_{e_1 \in E_{a_1}} \cdots \sum_{e_v \in E_{a_v}} P(e_1, \ldots, e_v \mid a_1 : \top, \ldots, a_v : \top, E) \\
&\qquad\qquad \epsilon(H \mid e_1, a_1 : \top, \ldots, e_v, a_v : \top, E)
\end{aligned}
\tag{11}
$$

In order to determine $Exp(\epsilon(H \mid E), a_1, \ldots, a_v)$, equation (11) can be simplified as follows:

$$
\begin{aligned}
&Exp(\epsilon(H \mid E), a_1, \ldots, a_v) \\
&= \sum_{e_1 \in E_{a_1}} \cdots \sum_{e_v \in E_{a_v}} \frac{P(e_1, \ldots, e_v, a_1 : \top, \ldots, a_v : \top, E)}{a_1 : \top, \ldots, a_v : \top, E} \\
&\qquad\qquad \epsilon(H \mid E \cup \{e_1, a_1 : \top, \ldots, e_v, a_v : \top\}) \\
&= \sum_{e_1 \in E_{a_1}} \cdots \sum_{e_v \in E_{a_v}} \left(\prod_{i=1}^{v} \frac{P(e_i, a_1 : \top, \ldots, a_v : \top, E)}{a_1 : \top, \ldots, a_v : \top, E} \right) \\
&\qquad\qquad \epsilon(H \mid E \cup \{e_1, a_1 : \top, \ldots, e_v, a_v : \top\}) \\
&= \sum_{e_1 \in E_{a_1}} \cdots \sum_{e_v \in E_{a_v}} \left(\prod_{i=1}^{v} P(e_i \mid a_1 : \top, \ldots, a_v : \top, E) \right) \\
&\qquad\qquad \epsilon(H \mid E \cup \{e_1, a_1 : \top, \ldots, e_v, a_v : \top\})
\end{aligned}
\tag{12}
$$

5.4.2 Multiple evidence sets

Certain investigating 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 investigating 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 $E_{a,1}, \ldots, E_{a,w}$:

$$
\begin{aligned}
Exp(&\epsilon(H \mid E), a) \\
&= \sum_{e_1 \in E_{a,1}} \cdots \sum_{e_w \in E_{a,w}} P(e_1, \ldots, e_w \mid a : \top, E) \\
&\qquad\qquad\qquad \epsilon(H \mid e_1, \ldots, e_w, a : \top, E) \\
&= \sum_{e_1 \in E_{a,1}} \cdots \sum_{e_w \in E_{a,w}} (\prod_{i=1}^{w} P(e_i \mid a : \top, E)) \\
&\qquad\qquad\qquad \epsilon(H \mid E \cup \{e_1, \ldots, e_w, a : \top\})
\end{aligned}
\tag{13}
$$

5.4.3 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 $H_1, \ldots, H_t$ by measuring entropy over these sets, given a set of pieces of evidence E, the following is required:

$$
\begin{aligned}
\epsilon(&H_1, \ldots, H_t \mid E) \\
&= -\sum_{h_1 \in H_1} \cdots \sum_{h_t \in H_t} P(h_1, \ldots, h_t \mid E) \log P(h_1, \ldots, h_t \mid E) \\
&= -\sum_{h_1 \in H_1} \cdots \sum_{h_t \in H_t} (\prod_{i=1}^{t} P(h_i \mid E)) \log(\prod_{i=1}^{t} P(h_i \mid E))
\end{aligned}
\tag{14}
$$

5.5 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.

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

Investigating action	EPE	NEER	REER
Knot analysis	0.98	26%	29%
Examine body	1.12	17%	19%
Search for cutting instrument	0.19	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 investigating actions

investigating action a (i.e. $\epsilon(H \mid E) - Exp(\epsilon(H \mid E), a)$) as a proportion of the maximal entropy under total lack of information. As such, it provides a means of assessing case progress:

$$NEER(H \mid E, a) = \frac{\epsilon(H \mid E) - Exp(\epsilon(H \mid E), a)}{\epsilon(H)} \quad (15)$$

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 investigating action possible:

$$REER(H \mid E, a) = \frac{\epsilon(H \mid E) - Exp(\epsilon(H \mid E), a)}{\epsilon(H \mid E)} \quad (16)$$

The results of these calculations are listed in Table 2 for the running example. This table presents the evaluation of a number of investigating actions after traces of anaesthetics have been discovered in johndoe's body. In particular, the second column displays the EPEs for investigating actions while the third and fourth columns show the corresponding NEER and REER values, respectively. Thus, the expected entropy reduction is the highest for knot analysis. In other words, knot analysis is likely to yield the most informative evidence and hence, it should be performed first (assuming that there are no other factors that need to be considered, such as the perishability of evidence).

To illustrate the computation, consider the investigating action $a = $ *knot analysis*. The current entropy, given that traces of anaesthetic have been found in johndoe's body (and before performing investigating action a), is

calculated as follows:

$$\epsilon(H \mid E) = -\,(P(\text{death of johndoe was homicide} \mid E_2 \cup \{\text{traces of anaesthetic in johndoe} : \top\}) \times$$

$$\log P(\text{death of johndoe was homicide} \mid E_2 \cup \{\text{traces of anaesthetic in johndoe} : \top\})) -$$

$$(P(\text{death of johndoe was suicide} \mid E_2 \cup \{\text{traces of anaesthetic in johndoe} : \top\}) \times$$

$$\log P(\text{death of johndoe was suicide} \mid E_2 \cup \{\text{traces of anaesthetic in johndoe} : \top\})) -$$

$$(P(\text{death of johndoe was accidental} \mid E_2 \cup \{\text{traces of anaesthetic in johndoe} : \top\}) \times$$

$$\log P(\text{death of johndoe was accidental} \mid E_2 \cup \{\text{traces of anaesthetic in johndoe} : \top\}))$$

$$= -\,0.40 \times \log 0.40 - 0.49 \times \log 0.49 - 0.11 \times \log 0.11 = 1.38$$

In addition, as previously shown, the entropy in the absence of any evidence, other than the hanging dead body of johndoe equals $\epsilon(H) = \epsilon(H \mid E_1) = 1.53$. The expected posterior entropy after performing a (and after the detection of traces of anaesthetics) can be computed in the same way as the illustrative example given in Section 5.3. This results in

$$Exp(\epsilon(H \mid E_2), a) = 0.98$$

From this, the NEER and REER can be calculated such that:

$$\text{NEER}(H|E_2, a) = \frac{1.38 - 0.98}{1.53} = 0.26$$

and

$$\text{REER}(H|E_2, a) = \frac{1.38 - 0.98}{1.38} = 0.29$$

6 Conclusions and Future Work

This paper has introduced a novel approach to generate plausible crime scenarios that are each sufficient to explain a given set of evidence. Such a scenario space is represented as a causal network of plausible states and events, factually and potentially available evidence, and investigating actions and hypotheses. The work offers an efficient means of storing a wide range of scenarios as the elements that different scenarios have in common need only be recorded once.

The paper has also extended Bayesian network-based techniques to identify, and further to refine, the most likely hypotheses by proposing effective evidence collection strategies. A set of maximum expected entropy reduction techniques have been devised that can determine the investigating actions that are expected to yield the most conclusive evidence. As such, this research enables the development of an intelligent decision support system for aiding police investigators in efficiently deciding on effective investigating actions. Clearly, this reflects the intended application of the present work – to assist the investigators to make, under uncertainty, decisions on how to reduce this uncertainty and therefrom, on how to optimise their "information acquisition strategy".

It is worth reiterating that the system described here is not intended to assist crime investigators in deciding whether or not the evidence against a suspect is sufficient, but how to find that evidence in the first place. This is categorically different from that faced by evidence evaluation in court. There, all relevant investigating actions have typically been carried out, and the issue is to decide on the basis of the results of these actions whether they prove the defendant's guilt to the required standards. However, the capability to consider more plausible scenarios and to gather the most likely useful evidence with limited resources will help deliver convincing evidence more timeously. This should ultimately help to reduce the likelihood of miscarriages of justice in court.

While the proposed system is promising, further improvements are required. As the probability distributions in the scenario fragments refer to prespecified assessments of the likely outcomes, which are typically subjectively given by expert investigators, they will often be described only generally. Thus, 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 [5, 9, 27] and linguistic probability theory [11]. Investigation into the use of symbolic probabilities therefore forms an interesting immediate future work.

Another piece of important further work concerns the relaxation of the important presumption made within this research that probability distributions governing the outcomes of different causal influences (and hence represented in distinct scenario fragments) which affect the same variable must be independent. Fortunately, the knowledge representation scheme adopted herein allows for this. In fact, information on the correlation amongst causal influences, specified by certain scenario fragments, could be added to the knowledge base, thereby explicitly representing how influences are interdependent. Implementation of this, however, requires considerable further studies.

Acknowledgements

Appendix: Notations

Symbol	Meaning
$\top$	true
$\bot$	false
$\wedge$	conjunction ("and" operator)
$\vee$	disjunction ("or" operator)
$\sum$	sum
$\prod$	product
x_i	a variable
a_i	an antecedent, assumption or action variable
n_i	a non-assumption variable
D_{x_i}	the domain of variable x_i (i.e. the set of all values that can be assigned to x_i)
$f : D_x \mapsto [0,1]$	a function f that maps values from the domain of x to the interval $[0,1]$.
$P(x : v)$	the probability that x is assigned value v
$P(x_1 : v_1 \mid x_2 : v_2)$	the probability that x_1 is assigned value v_1 given that it is known that x_2 is assigned value v_2
v_i	a value from a variable's domain
c_i	a value from a consequent variable's domain
$x : v$	an assignment of value v to variable x
A	a set of assignments
SF_i	a scenario fragment
$\oplus$	combination operator of a scenario fragment
$A \xrightarrow{SF_i} x : c$	scenario fragment SF_i assigns value c to x given the set of assignments A
H	a set of hypotheses
$\epsilon(H)$	entropy over H
E	a set of pieces of evidence
$\epsilon(H \mid E)$	entropy over H given E
$Exp(\epsilon(H \mid E), a_1, \dots, a_v)$	expected entropy over H given E after performing investigating actions $a_1, \dots, a_v$

References

[1] C.G.C. Aitken and F. Taroni. *Statistics and the evaluation of evidence for forensic scientists, 2nd edition*. Wiley, 2004.

[2] C.G.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] 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.

[4] 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.

[5] 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.

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

[7] D. Dixon. Police investigative procedures: Changing legal and politic contexts of policing practices. In C. Walker and K. Stramer, editors, *Miscarriages of justice*, pages 65–82. Blackstone, London, 1999.

[8] 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.

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

[10] S. Greer. Miscarriages of criminal justice reconsidered. *Modern Law Review*, pages 58–74, 1994.

[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] J.B. Kadane and D.A. Schum. *A Probabilistic Analysis of the Sacco and Vanzetti Evidence*. John Wiley and Sons, New York, 1996.

[14] M. McConville. Weaknesses in the british judicial system. In *Times Higher Education Supplement 3.11*. 1989.

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

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

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

[18] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2000.

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

[20] M. Redmayne. The royal commission and the forensic science market. In M. McConville and L. Bridges, editors, *Criminal Justice in Crisis*, pages 219–228. Aldershot: Edward Elgar, 1994.

[21] P. Roberts. What price a free market in forensic science services? the organization and regulation of science in the criminal process. *British Journal of Criminology*, 36:37–60, 1996.

[22] B. Schafer and J. Keppens. Knowledge based crime scenario modelling. *Expert Systems with Applications*, 30(2):203–222, 2006.

[23] D. Schum. *The Evidential Foundations of Probabilistic Reasoning*. Wiley, 1994.

[24] F. Taroni, C.G.C. Aitken, P. Garbolino, and A. Biedermann. *Bayesian Networks and Probabilistic Inference in Forensic Science*. Wiley, 2006.

[25] V. Torra. Aggregation operators and models. *Fuzzy Sets ann Systems*, 156(3):407–410, 2005.

[26] C. Walker and K Starmer. *Justice in Error*. Blackstone, London, 1993.

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