On the Role of Model-based Reasoning in Decision Support in Crime Investigation

Jeroen Keppens and John Zeleznikow
The University of Edinburgh
Centre for Forensic Statistics and Legal Reasoning
jeroen@dai.ed.ac.uk, jzelezni@staffmail.ed.ac.uk

ABSTRACT

Crime investigation is a complex task involving vast amounts of information and requiring many different types of expert knowledge. Crime investigators would therefore benefit from the use of decision support systems to help manage this information and to provide knowledge to help solve the more complex problems. Current research efforts in this area have focussed on the information management side of the problem and tend to steer clear of formalising expert knowledge. This is understandable since conventional knowledge based systems lack the robustness needed to cope with the variety of circumstances that can be encountered during criminal investigation. However, similar problems have been encountered in the physical systems domain and were tackled by means of novel model-based reasoning techniques. This paper explores the use of robust model-based reasoning approaches to model expert knowledge for crime investigation and it presents a framework for such systems. The preliminary ideas presented in this paper are illustrated by means of practical examples produced during ongoing work in the development of a system for differentiating between homicidal, suicidal, accidental and natural deaths.

KEY WORDS

decision support system, crime investigation, artificial intelligence, model-based reasoning

1 Introduction

The task of crime investigators is to discover crimes, their perpetrators and sufficient evidence to enable a prosecution whilst being economical with scarce resources and avoid harming innocent people. Because potentially contradictory goals need to be achieved and because complex, often specialist, information is being generated in the process, crime investigations are difficult to manage. To aid in the management of such complex tasks, decision support systems are often proposed.

Many different approaches have been used in the past to design and implement decision support systems. Conventional rule-based systems are good at replicating the heuristics experts use. Case based reasoning systems may recall similar, previously solved cases and adapt their solutions to the current requirements. Planning and scheduling systems can assist in the operational management of a case. Despite their differences, all of these approaches share the requirement that the content of all problem instances can be generalised and represented formally. But crime investigation problems are too varied for this to be achievable.

The main questions to be addressed in crime investigation are: (i) what hypotheses are best supported by the available evidence, and (ii) what potential sources of evidence best distinguish between competing hypotheses. These questions are similar to those answered by diagnostic systems. However, they are particularly difficult in crime investigation. Even though the set of component elements (e.g. types of crime, types of evidence, etc.) of crime investigation problems are by themselves are relatively limited, they can occur in an almost infinite number of different combinations. Since it is impossible to foresee all of these scenarios in advance, an approach is needed to deal with the issue of robustness of the decision support system.

Model-based reasoners have been devised to deal with this issue of robustness (or the portability over different problem scenarios) [8]. Instead of employing problem specific knowledge directly to solve a problem instance, they employ domain knowledge to create models of the problem at hand and apply problem independent techniques onto those models. In this way, model-based reasoners can handle a wider variety of situations and hence, they are more robust. This paper presents an analysis of how such model-based approaches can be applied to provide decision support for crime investigation. Future work will build on these ideas through the design of inference mechanisms that implement them.

2 Model-based Diagnosis

The aim of a diagnostic system is to determine the hypothesised state of a system via a set of observable consequences or symptoms of that state. Conventional rule-based approaches perform this task by using expert knowledge to relate observable pieces of evidence to plausible hypotheses. In other words, such diagnostic systems simulate by means of deduction what is essentially an abductive task. Therefore, rule-based diagnosers are only practical in domains where the potential problems that can be encountered are well understood. Many domains, including crime investigation, do not fall into this category.

As opposed to conventional approaches, a model-based diagnostic system employs a knowledge based system K that can determine for a given system S whether it satisfies a hypothesis h (i.e. $S, K \vdash h$) or whether a symptom e logically follows from it (i.e. $S, K \vdash e$). Using K, the diagnostic engine can determine which hypotheses in a set H are consistent with the set of available symptoms E by searching:

$$\{h \in H \mid \exists S, (\forall e \in E, (S, K \vdash e)) \wedge (S, K \vdash h)\}$$

Two main approaches are used to derive the knowledge base K:

- *Constraint propagation/satisfaction* can be applied to a mathematical model to discover the sets of conditions with which certain sets of hypotheses and symptoms are consistent. For example, these techniques are used to determine which hypothesised faulty operating conditions generate the undesirable symptoms in an engineering system [2].
- *Automated modellers* are used to generate mathematical, or other, models describing hypothesised configurations of the diagnosed system, and which may produce the observed behaviour. For example, automated modelling techniques have been used to find which plausible configurations of an eco-system may be responsible for observed disasters [7].

The constraint propagation/satisfaction algorithm and/or the automated modeller generate two types of knowledge: (1) inferences describing which symptoms and/or hypotheses can be derived from which partial models, and (2) which partial models describe an inconsistency. In model-based diagnosis, this knowledge is stored in a truth maintenance system (TMS), usually an assumption-based truth maintenance system (ATMS).

In general, an ATMS [3] assists a problem solver by maintaining which consistent sets of uncertain assumptions justify which propositions. For this purpose, it takes a set of assumptions (problem solver datums whose truth needs to be established), a set of nodes (other problem solver datums involved in the inferences) and inferences of the form:

$$n_1 \wedge \ldots n_k \rightarrow n_c$$

where $n_1 \ldots n_k$ correspond to assumptions or nodes and n_c represents a node or falsehood (if $n_c = \perp$). Based on this information, the ATMS can compute a label $\mathcal{L}(n)$ for each node n. A label $\mathcal{L}(n)$ describes a disjunctive normal form expression, $\vee_{E_i \in \mathcal{L}(n)} \wedge_{a_{ij} \in E_i} a_{ij}$ that is

- *consistent*: no conjunction $\wedge_{a_{ij} \in E_i} a_{ij}$ support falsehood,
- *minimal*: there is no pair of sets $E_1, E_2 \in \mathcal{L}(n)$ such that $E_1 \subset E_2$,
- *complete*: for every conjunction of assumptions E from which the node n logically follows, a superset of E can be found in $\mathcal{L}(n)$,
- *sound*: the node n logically follows from every conjunction of assumptions E in $\mathcal{L}(n)$.

This TMS contains the knowledge base K and it provides an efficient way to determine from which set of consistent scenarios a particular piece of evidence or hypothesis can be derived.

3 Model-based Crime Investigation

For the process of crime investigation, the diagnostic tasks are as follows. The system under diagnosis is a space of plausible scenarios that is presumed to include unlawful actions, possibly a crime. The symptoms are the pieces of evidence that are or can be collected by the investigators. And finally, the hypotheses about the scenario, which must be determined by the investigators, relate to important properties of the crime, such as the type of the crime, the perpetrator, etc.

A representative example of such a diagnosis task is discrimination between homicidal, suicidal, accidental and natural death. Here, the scenario is a sequence of events leading up to a dead body. The evidence (or symptoms) can vary widely but include features of the dead body, DNA, trace evidence, CCTV footage and witness statements. The hypotheses are “homicide”, “suicide”, “accident” or “natural causes”. This problem presents itself whenever police officers encounter a dead body. It is often a non-trivial crime investigation task because the evidence that enables the differentiation is often very subtle and possibly even hidden. It is further complicated by the fact that police officers often have little experience in making this assessment, because in most countries, death by homicide or suicide is far rarer than death by natural causes.

For the human crime investigator, the most difficult aspects of this diagnosis task involve collecting and organising a representative body of evidence and taking into consideration a sufficiently large set of hypotheses. Miscarriages of justice commonly occur when the investigation focuses on a single hypothesis, or even a single corresponding scenario, and evidence is predominantly collected to support only this hypothesis. To help avoid such mistakes and the associated human and financial costs, a decision support system for crime investigation can aid by considering a set of plausible scenarios supporting alternative hypotheses whilst efficiently guiding the evidence collection to discriminate between multiple hypotheses. To that end, a decision support system for crime investigation requires the following functionalities:

- *Constructing scenarios that support the body of evidence*. The crime investigation equivalent of models matching hypothesised behaviours is scenarios describing sequences of events. Such scenarios can be composed from recurring (and therefore reusable) component events as well as the pieces of evidence (i.e. the symptoms) each event is likely to produce.
- *Ranking the scenarios in terms of plausibility*. Because scenarios are derived abductively from their observed consequences, the decision support system will normally be able to generate multiple plausible scenarios.

Yet, not all scenarios are equally plausible and, therefore, uncertainty reasoning techniques should be employed to make these distinctions.

- *Seeking out discriminatory pieces of evidence*. Similar to model-based diagnosis, crime investigation is an iterative task: the available evidence is interpreted, preliminary conclusions as to the plausibility of the hypotheses are drawn and, based on this analysis, new evidence is searched to reach a more refined interpretation of the evidence. By analysing how evidence stems from plausible scenarios and how hypotheses are inferred from these scenarios, model-based diagnostic techniques can be employed to seek out what evidence would best discriminate between the available hypotheses.

The remainder of this section will discuss each of these features in more detail.

3.1 Scenario composition

Any approach to apply model-based reasoning techniques to decision support systems for crime investigations is crucially dependent on the development of models of the problem at hand. Such models describe how evidence and hypotheses are inferred and they should encompass the entire range of plausible scenarios. One of the most widely used approaches to generate such models, in areas such as physics, engineering, botany and ecology, is compositional modelling [5].

A compositional modeller derives a model from a knowledge base of component models. Because the component models are parts of scenarios, they are called scenario fragments. Each scenario fragment describes how the state of the world can change through an event. To that end, a scenario fragment consists of a set of assumptions, a set of prerequisite states, a set of consequent states and an event. A scenario fragment formalises the knowledge that under the set of assumptions, the prerequisite states imply that the event occurs and the consequent states are become part of the world model.

Definition 1 (Scenario fragment) *A scenario fragment is a tuple $\langle A, S^p, S^c \rangle$ where A is a set of assumptions, S^p is a set of prerequisite states and events, and S^c is a set of consequent states and events, such that*

$$\wedge_{a \in A} a \rightarrow (\wedge_{s^p \in S^p} s \rightarrow \wedge_{s^c \in S^c} s^c)$$

The following examples may clarify these concepts:

$$\text{ability(self-defence}(V))\wedge \text{assault}(P,V) \rightarrow \text{self-defence}(V,P) \quad (1)$$

$$\text{lethal-beating}(P,V) \rightarrow \text{homicide}(V) \quad (2)$$

Scenario fragment (1) states that if a victim V is being assaulted by a perpetrator P, and (s)he is capable of self-defence, then V will defend against and assault by P. Scenario fragment (2) describes that a lethal beating of V by P implies that the death of V is a homicide case.

Not every combination of states will be possible in the real world. To prevent the decision support system from considering such inconsistent combinations of states, a special state, called the inconsistent state, $\perp$ is introduced. Certain scenario fragments may contain $\perp$ as a consequent state, indicating that the combination of assumptions and prerequisite states it contains is inconsistent. For example, presuming that lethal-beating(A, B) denotes the event that A beats B, causing the death of B, then (lethal-beating$(P, V) \wedge$ lethal-beating(V, V)), with $P \neq V$ is such an inconsistency[1].

In order to find the hypotheses that are consistent with the available evidence (and, as discussed in section 3.2, how plausible they are), a sound and complete inference mechanism is required to compute those scenarios S from which the available evidence E logically follows in the knowledge base K (i.e. $S, K \vdash E$). An inference mechanism is deemed *sound* if it only produces consistent scenarios supporting the all the evidence. It is considered *complete*, if it can find all the scenarios in the knowledge base that support the available evidence.

A suitable complete and sound inference mechanism is the ATMS. In general, the knowledge contained in the scenario fragments can be entered into an ATMS as follows. By means of a backward chaining search, all scenario components supporting the available evidence can be determined in the knowledge base. A conventional forward chaining search can determine inconsistencies, hypotheses and potential, but yet unexplored pieces of evidence. Once these inference procedures have been completed, sets of assumptions supporting plausible scenarios can be computed.

Let n be node inferred as $(\wedge_{e \in E}) \wedge h \rightarrow n$, where h is a single hypothesis and E is the set of available evidence. Then, each set $A \in \mathcal{L}(n)$ is a set of assumptions from which a plausible scenario follows that contains the hypothesis h as well as the available evidence E. An example may clarify these concepts.

Assume a scenario in which the dead body of a victim V is found. The pathology report indicates that V died from injuries caused by beating and that the victim has irregular fingernails. The hypotheses set is $H = \{\text{homicide}(V), \text{suicide}(V)\}$. Figure 1 shows a sample scenario, which can be generated by means of the automated modelling approach discussed herein. This scenario describes how V is assaulted by a perpetrator P, which eventually causes V's death and V defends him/herself (causing the irregular fingernails). This scenario is derived from the set of assumptions: {moves-to(V, L), moves-to(P, L), lethal-beating(P, V), ability(self-defence(V)), includes(defence(V, P), scratching(V, P))}. This set of assumptions describes as set of possible, but unknown, states and events causing the scenario.

Also shown in figure 1 is part of a scenario that defeats the homicide hypothesis: V commits suicide through self-beating. However, the latter scenario is expected to generate certain pieces of evidence that may or may not be found.

[1] Although multiple beatings may *contribute* to the victims death, at most one can be the direct *cause*.

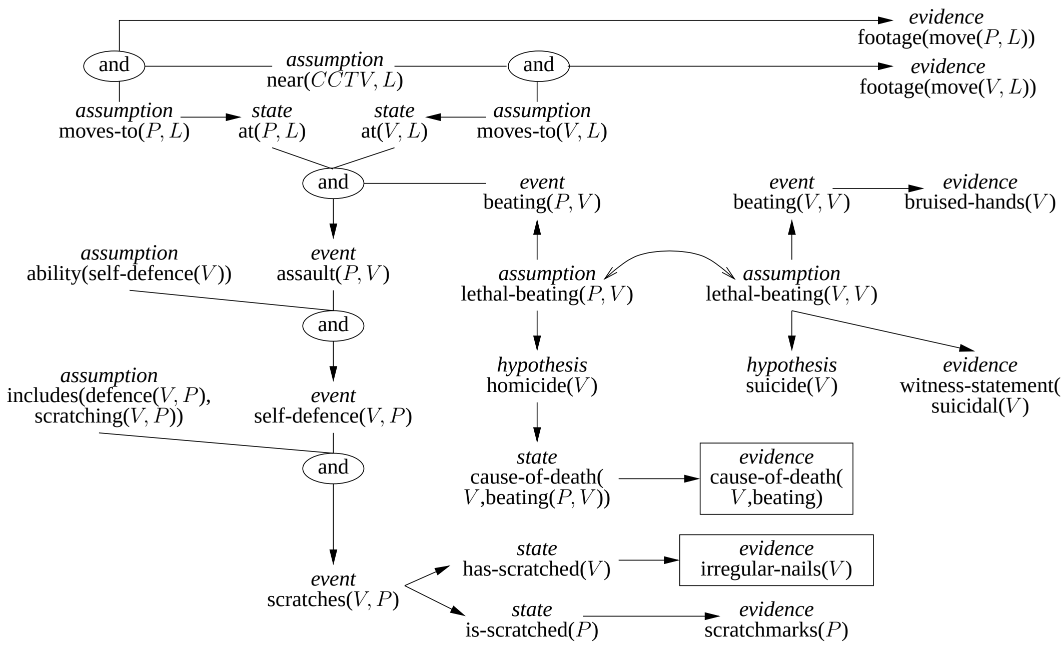

Figure 1. Sample scenario and related evidence and hypothesis

3.2 Ranking scenarios

Section 3.1 explained how a truth maintenance system can be populated with a space of plausible scenarios describing alternative hypotheses. Each of these scenarios presupposes a number of assumptions. The assumptions describe unknown events. Incomplete knowledge as to whether these assumptions are correct gives rise to uncertainty about whether or not a scenario has occurred.

3.2.1 Prior likelihoods

Let $\wedge_{a_{ij} \in A_i}$ be a minimal and consistent set of assumptions from which the entire set of evidence E and a single hypothesis h logically follows: $\wedge_{a_{ij} \in A_i} a_{ij}, K \vdash ((\wedge_{e \in E} e) \wedge h)$. Then, the prior probability that corresponding scenario occurs, given K can be expressed as $P(\wedge_{a_{ij} \in A_i} a_{ij}, K \vdash ((\wedge_{e \in E} e) \wedge h))$.

Calculating this probability is almost impossible. In general, limited information is available to determine likelihood of the assumptions and the inferences based on them. For example, the prior probability that people injure themselves is very hard to quantify. For the purpose of the type decision support envisioned in this paper, it is best to keep an open mind. It should therefore be sufficient to provide a partial ordering over alternative scenarios. Therefore, some simplifying assumptions will be made. Firstly, it is presumed that no uncertainty arises from the logical inference. Hence

$$P(\wedge_{a_{ij} \in A_i} a_{ij}, K \vdash ((\wedge_{e \in E} e) \wedge h)) = P(\wedge_{a_{ij} \in A_i} a_{ij})$$

Secondly, it is presumed that important distinctions between prior probabilities can be represented as partial orderings. In crime investigation, the most important way of establishing such an ordering is by deeming closed-world assumptions more probable than their defeaters.

Many of the assumptions used in the scenario descriptions are so-called closed-world assumptions. Closed-world assumptions are presumptions that enable the model construction algorithm to ignore uncommon complications. As such, closed-world assumptions enable default reasoning without entirely ignoring defeaters of the default. An example of a closed-world assumption is the presumption that the victim of an assault is capable of some form self-defence. This is an important presumption in the investigation since it implies that an assault on the victim can diagnosed by means of defensive injuries (e.g. irregular fingernails in figure 1). However, the absence of defensive injuries does not exclude the assault scenario if one of the defeaters is presumed. For example,

$$\text{restrained}(V) \rightarrow \neg\text{ability}(\text{self-defence}(V))$$

states that if V is restrained, the aforementioned closed-world assumption and its consequences is inconsistent.

Because ability(self-defence(V)) is the default assumptions (in the absence of other information), it is appropriate to presume that $P(\text{ability}(\text{self-defence}(V))) \leq P(\neg\text{ability}(\text{self-defence}(V)))$.

3.2.2 Posterior likelihoods

By collecting more information in the form of evidence, the probabilities can be better differentiated from one another. The conventional approach to compute conditional probabilities $P(\wedge_{a_{ij} \in A_i} a_{ij} \mid E)$, where E denotes the set of available evidence, is by means of Bayes law:

$$P(\wedge_{a_{ij}\in A_i} a_{ij} \mid E) = \frac{P(E \mid \wedge_{a_{ij}\in A_i} a_{ij}) P(\wedge_{a_{ij}\in A_i} a_{ij})}{P(E)} \tag{3}$$

As with the prior probabilities, it is extremely difficult to obtain good estimates for all the conditional probabilities involved in (3). For some types of evidence, such as DNA and trace evidence, good statistics are available for both $P(E \mid \wedge_{a_{ij}\in A_i} a_{ij})$ and $P(E)$. For example, research efforts into the distribution, transfer and shedding of fibres has led to a body of statistics (i) on the likelihood of encountering given types of fibre in garments (i.e. $P(e)$ where e represents the discovery of certain fibres) and (ii) on the likelihood of encountering fibres after presumed types of contact between garments, which may have led to fibre transfer (i.e. $P(e \mid \wedge_{a\in A} a)$ where $\wedge_{a\in A} a$ describes the presumed events) [6].

Yet, in general, it is difficult to produce such statistics. The probabilities $P(e)$ and $P(e \mid \wedge_{a\in A} a)$ are hard to quantify in the case of many types of evidence. For some types of evidence, such as witness statements, the important influencing factors, such as the reliability and objectivity of witnesses, the exact conditions in which the observation was made, and how well the witness remembers the events, are difficult to measure. The interpretation of some other types of evidence, such as CCTV footage and fingerprints, is done by human experts who may not be able to produce accurate statistics on the reliability of their work.

One way of dealing with this problem involves the use of interval calculus [1] to compute ranges of probabilities covering the entire span of plausible outcomes. Another approach of expressing such crude probabilities is discussed in [4]. This work suggests the use of different classes of probability, one class being an order of magnitude greater than the next. In this way, comparing products of probabilities is reduced to counting the number probabilities of each class and the result is dominated by distinctions between the more reliable types of evidence.

In most cases, it is only practical to elicit probabilities over subsets of assumptions and corresponding pieces of evidence. Therefore, some independence assumption must be made herein, as in any other diagnostic system. Let $e \Rightarrow A$ denote that the likelihood of a set of assumptions A is affected by a piece of evidence e, or $P(\wedge_{a\in A} a \mid e) \neq P(\wedge_{a\in A} a)$, and let $E(A)$ be the set of all pieces of evidence that affect the likelihood of a set of assumptions, or $E(A) = \{e \mid e \Rightarrow A\}$. Now, the simplifying presumption can be made that if two assumptions have different sets of pieces of evidence that potentially affect them, their prior and posterior probabilities are independent from one another. More formally, if $E(\{a_i\}) \neq E(\{a_j\})$, then $P(a_i \wedge a_j) = P(a_i)P(a_j)$ and $P(a_i \wedge a_j \mid E) = P(a_i \mid E)P(a_j \mid E)$. By partitioning the set of assumptions A_i underlying a plausible scenario into a set $\{A_i^1, \ldots, A_i^k\}$, the likelihood of that scenario can be computed as:

$$P(A_i \mid E) = \prod_{A_i^j} P(\wedge_{a\in A_i^j} a \mid E(A_i^j))$$

Based on these notions, a partial ordering of scenarios can be established. Ordering two scenarios involves comparing the prior probabilities of the assumptions for which no posterior probabilities can yet be computed (i.e. $E(A) = \emptyset$) and comparing the combined posterior probabilities of all other assumptions. Formally, a scenario supported by A_1 no more plausible than a scenario supported by A_2 if:

$$\left[P(\wedge_{a\in\{a\in A_1 \mid E(\{a\})=\emptyset\}}) \leq P(\wedge_{a\in\{a\in A_1 \mid E(\{a\})=\emptyset\}})\right] \wedge$$
$$\left[\prod_{A_1'\subseteq A_1, A_2'\subseteq A_2, E(A_1')=E(A_2')} P(\wedge_{a\in A_1'} a \mid E(A_1')) \leq \right.$$
$$\left.\prod_{A_1'\subseteq A_1, A_2'\subseteq A_2, E(A_1')=E(A_2')} P(\wedge_{a\in A_2'} a \mid E(A_2'))\right]$$

3.3 Discriminatory evidence

By means of the ideas presented in section 3.2, sets of plausible scenarios can be ranked with respect to one another. Such rankings can not only be used to determine which scenarios are most likely to describe the events that lead up to the available evidence, but also to establish to guide the evidence collection strategy.

To support the collection of evidence, diagnostic systems aim to reduce to the entropy (or chaos) in the existing set of plausible hypotheses. In better understood domains, the entropy $\epsilon(E)$ over the available hypotheses given a set of evidence E is computed by:

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

But, as argued before, the probabilities in (4) are very hard to quantify in general. By using the partial ordering of scenario, entropy over the plausible scenarios supporting different hypotheses can be computed, however. Presume that some heuristic is used to reduce the space of plausible scenario to the most likely ones, and let $c(h \mid E)$ be the total number scenarios that are supports hypothesis h. Then, then the current entropy can be computed as:

$$\epsilon(E) = \sum_{h\in H} c(h \mid E) \ln c(h \mid E) \tag{5}$$

An optimal decision making strategy is one that reduces the expected future entropy as much as possible. Let x_i be a type of evidence that can be collected, and let D_i be the set of possible outcomes of pursuing a search for x_i. For example, if x_i corresponds to searching for fingerprints at some location, then D_i might contain different sets of types of fingerprint found (e.g. the victim's fingerprint, an unknown fingerprint found, etc.) as well as nothing found. The expected entropy $\epsilon(x_i)$, after searching for evidence type x_i, can be computed as follows:

$$\epsilon(x_i) = \sum_{v_{ij}\in D_i} P(x_i : v_{ij}) \epsilon(E \cup \{x_i : v_{ij}\}) \tag{6}$$

where $P(x_i : v_{ij})$ is the probability of outcome v_{ij} and $\epsilon(E \cup \{x_i : v_{ij}\})$ is the new entropy assuming that the outcome would be v_{ij}. Given $\epsilon(x_i)$ for all possible evidence collection decisions, the optimal evidence collection strategy selects the decision with the lowest expected entropy.

Note that, once again, some crude estimate of a probability is required. If there is no way of telling in advance what $P(x_i : v_{ij})$ is likely to be, $P(x_i : v_{ij})$ can always be presume to be equal to $\frac{1}{|D_i|}$, where $\mid D_i \mid$ denotes the cardinality of D_i. In other case, e.g. when the room where fingerprints are to be collected might have been cleaned after the crime, rough estimates of the likelihoods of plausible outcomes are obvious.

4 Conclusions and Future Work

The work presented in this paper is a first step in the development of decision support systems that confront crime investigators. More specifically, the paper has argued for the use of model-based reasoning techniques to tackle this problem and established their role. By its nature, the task of crime investigation does not lend itself to be captured by a set of rules. The overall task of determining which potentially unlawful events have occurred requires consideration of unforseeable combinations of such events, and model-based reasoning techniques have been devised with this challenge in mind.

Overall, three potential tasks for model-based diagnosis have been identified and explored in this paper. In the first instance, automated model construction approaches, such as compositional modelling, can be used to construct plausible scenarios that explain the available evidence as well as a possible hypothesis. Once a complete space of scenarios has been constructed, it can be employed for further analysis. In particular, the dependency of scenarios upon sets of closed-world assumptions can be exploited to rank scenarios in terms of their relative likelihoods. The partial ordering imposed, in this way, over the scenario space can, in turn, be used to compute an optimal evidence collection strategy.

The main drawback of the approach proposed herein is the knowledge acquisition bottleneck that is commonly encountered in the deployment of knowledge intensive systems. Indeed, any implementation will require a major knowledge acquisition effort to generate complete and valid scenario fragments. The effect of this issue can be somewhat limited by restricting the scope of the domain. As such, it is anticipated that this work will be most useful in the construction of specialist tools for the analysis of certain types of major crimes. Currently, a prototype is under development that aims to distinguish between homicidal, suicidal, accidental and natural deaths.

Aside from these implementational issues, a number of important theoretical issues remain to be addressed by future research. In particular, an approach to compute with crude probabilities, suitable for crime investigation, must be formalised. This paper has identified two potential approaches, but it is not yet clear which one is most suitable for expressing the type of uncertainty that crime investigators are confronted with. The development of the prototype knowledge base for the aforementioned application domain is expected to improve our understanding of this problem.

Finally, the incorporation of an appropriate framework for temporal reasoning (and to a lesser extent, for spatial reasoning) is necessary to refine the proposed scenario composition. The combination of events contained in a scenario may be constrained by time durations in addition their relative order. In a small number of cases, these time durations may be very important. Consider, for example, the case where a person dies due to a poisoning process that must have taken several weeks or months. In order to differentiate between homicide or accidental death in this case, it is important to know whether significant other events have occurred throughout or at a much smaller time intervals. Again, these issues must be addressed in future work.

Acknowledgements

The authors would like to thank (in alphabetical order) Michael Bromby, Anthony Forde, John Kingston, Patricia McKellar, Moira MacMillan and Burkhard Schafer for useful discussions.

References

[1] Alefeld, G. and Herzberger, J. *Introduction to Interval Computation*. Academic Press, 1983.

[2] Davis, R. and Hamscher, W. Model-based reasoning: troubleshooting. In Shrobe, H., editor, *Exploring Artificial Intelligence*, pages 297–346. Morgan-Kaufmann, 1988.

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

[4] de Kleer, J. Using crude probability estimates to guide diagnosis. *Artificial Intelligence*, 45:381–391, 1990.

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

[6] Robertson, J. and Grieve, M., editors. *Forensic Examination of Fibres*. Forensic Science Series. Taylor & Francis, 1999.

[7] Struss, P. Artificial intelligence for nature - why knowledge representation and problem solving should play a key role in environmental decision support. In Haasis, H.D. and Ranze, K.C., editors, *Computer Science for the Environmental Protection '98*. Metropolis Verlag, 1998.

[8] Zeleznikow, J, and Hunter, D. *Building Intelligent Legal Information Systems: Representing and Reasoning in Law*. Computer/Law Series. Kluwer, 1994.