

Chapter 22

REASONING ABOUT EVIDENCE USING BAYESIAN NETWORKS

Michael Kwan, Kam-Pui Chow, Frank Law and Pierre Lai

Abstract There is an escalating perception in some quarters that the conclusions drawn from digital evidence are the subjective views of individuals and have limited scientific justification. This paper attempts to address this problem by presenting a formal model for reasoning about digital evidence. A Bayesian network is used to quantify the evidential strengths of hypotheses and, thus, enhance the reliability and traceability of the results produced by digital forensic investigations. The validity of the model is tested using a real court case. The test uses objective probability assignments obtained by aggregating the responses of experienced law enforcement agents and analysts. The results confirmed the guilty verdict in the court case with a probability value of 92.7%.

Keywords: Digital evidence, hypotheses, probability, Bayesian networks

1. Introduction

Like other forensic disciplines, digital forensics involves the formulation of hypotheses based on the available evidence and facts, and the assessment of the likelihood that they support or refute the hypotheses. Although substantial research has focused on principles and tools for retrieving digital evidence [7, 8, 10], little, if any, work has examined the accuracy of hypotheses based on the evidence.

Without reliable and scientific models, the conclusions made by digital forensic analysts can be challenged on the grounds that they are mere speculation. The problem is exacerbated by the fact that forensic conclusions derived from the same digital evidence can vary from analyst to analyst. This can severely impact the reliability of digital forensic findings as well as the credibility of analysts. Speculation and subjective

views offered by forensic analysts under the guise of expert opinion have little (if any) value in legal proceedings [6].

This paper presents a formal model for reasoning about digital evidence. The model, which is based on probability distributions of hypotheses in a Bayesian network, quantifies the evidential strengths of the hypotheses and, thereby, enhances the reliability and traceability of the analytical results produced by digital forensic investigations. The validity of the model is investigated using a real court case involving the illegal dissemination of a movie using the BitTorrent peer-to-peer network.

2. Background

Forensics is the process of analyzing and interpreting evidence to determine the likelihood that a crime occurred. Many researchers (see, e.g., [4, 12, 15, 19, 20]) argue that that this process should cover the formulation of hypotheses from evidence and the evaluation of the likelihood of the hypotheses for the purpose of legal proceedings.

Aitken and Taroni [1] state that likelihood is an exercise in hypothetical reasoning. It denotes the degree of belief in the truth of a hypothesis. In the scientific community, belief is often expressed in terms of probability. Probability theories provide mechanisms for deducing the likelihood of hypotheses from assumptions. Although probabilistic methods may be useful for proving or refuting the hypotheses involved in a criminal investigation, Jones and co-workers [9] argue that obtaining all the probability distributions for the entailing evidence is impractical. Given the large volume of evidence involved, it is not feasible to obtain the joint probability distributions for all possible evidential variables. Moreover, simple probabilistic methods do not capture the complex dependencies that exist between items of evidence; therefore, the methods have limited value from an analytical point of view [5]. Indeed, many researchers [5, 11, 16] emphasize that comprehensive probabilistic models should accurately model the conditional dependencies existing between items of evidence.

A criminal investigation is an abductive diagnosis problem [16]. However, it is difficult to design a model that can deterministically describe all the assumptions involved in an investigation. Poole [18] has attempted to address this issue by proposing a model that describes crime scenarios non-deterministically using symbolic logic and probabilistic Bayesian methods. Unfortunately, Poole's model is too abstract to be applied in real scenarios.

It is important to observe that digital events are discrete computer events that are deterministic in nature and have a temporal causal sequence. Therefore, it is common practice for digital forensic analysts to establish their abductive reasoning based on the existence or validity of the causal events that entail their hypotheses. However, it is difficult to have consistent models that determine the supporting events for hypotheses. Different analysts may attach different events to the same hypothesis. Even if they agree on the same set of events, they usually assign different (subjective) probabilities to the events.

Analysts also must reason about hypotheses in the face of missing and/or uncertain information about events. The events for which evidence is available may not prove the complete truth of the hypotheses; however, they can be used very effectively to compute degrees of likelihood for the hypotheses. Consequently, probabilistic approaches are well suited to developing formal models for reasoning about digital evidence in criminal investigations.

3. Bayesian Networks

Before we discuss Bayesian networks, it is important to emphasize that digital evidence deals with “past” events that were caused by some other hypothetical events that have to be verified. For example, if a suspect had child pornography on his computer, he may have downloaded it from a pornographic web site, which could be verified by the presence of the URL in the history file of his browser.

A Bayesian network uses probability theory and graph theory to construct probabilistic inference and reasoning models. It is defined as a directed acyclic graph with nodes and arcs. Nodes represent variables, events or evidence. An arc between two nodes represents a conditional dependency between the nodes. Arcs are unidirectional and feedback loops are not permitted. Because of this feature, it is easy to identify the parent-child relationship or the probability dependency between two nodes.

A Bayesian network operates on conditional probability. For example, if the occurrence of some evidence E is dependent on a hypothesis H , the probability that both H and E occurred, $P(H, E)$, is given by:

$$P(H, E) = P(H)P(E|H). \quad (1)$$

According to the multiplication law of probability, which expresses commutativity, if H is relevant for E , then E must also be relevant for

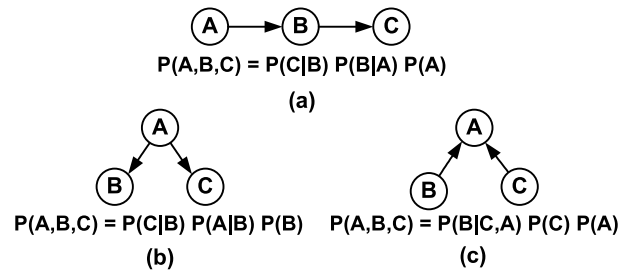


Figure 1. Bayesian network connections: (a) Serial; (b) Diverging; (c) Converging.

H. The corresponding joint probability expression is:

$$P(H, E) = P(H)P(E|H) = P(E)P(H|E), \quad (2)$$

and, hence,

$$P(E|H) = \frac{P(E)P(H|E)}{P(H)}. \quad (3)$$

Equation (3) is the celebrated Bayes' Theorem. From a statistical point of view, it denotes the conditional probability of *E* caused by *H*. This is also referred as the likelihood ratio of *H* given *E*. It denotes the degree of belief that *E* will occur given a situation where *H* is true.

$P(H|E)$ is the posterior probability, i.e., the probability that when *E* is detected *H* has actually occurred. $P(H)$ denotes the prior probability of *H* at a stage where the evidence is not yet presented. $P(E)$ is the prior probability of *E*, which is sometimes referred to as a normalizing constant. Therefore, the above expression can be formalized as:

$$\text{likelihood ratio} = \frac{\text{posterior probability} \times \text{normalizing constant}}{\text{hypothesis prior probability}}. \quad (4)$$

Since the likelihood ratio is proportional to the posterior probability, a larger posterior probability denotes a higher likelihood ratio. In the evidentiary context, it also means that the greater the evidence supporting the hypothesis, the more likely that the hypothesis is true.

A Bayesian network has three elementary connections between its nodes that represent three different types of probability distributions (Figure 1). For a serial connection, if *B*'s evidential state is unknown, then *A* and *C* are dependent on each other. In other words, there is an evidential influence between *A* and *C* if the evidential state of *B* is unknown. However, if *B*'s state is known, then *A* and *C* are independent of each other; this means that *A* and *C* are conditionally independent of each other given *B*. In a diverging connection, the same conditional

independence is observed for A and C , i.e., if B 's state is known, then A and C are independent. In a converging connection, if B 's state is unknown, then A and C are independent. In other words, unless the state of B is known, A and C can influence each other.

4. Proposed Model

A real case involving the distribution of a pirated movie via the BitTorrent peer-to-peer network is used to demonstrate the utility of the Bayesian network model. The digital evidence discussed in this paper was presented in court during the criminal trial.

4.1 The BitTorrent Case

The defendant in the case was alleged to have used his computer to distribute a pirated movie on the Internet using BitTorrent [13]. The defendant had the optical disk of the movie in his possession. He copied the movie from the optical disk to his computer and then used BitTorrent to create a "torrent file" from the movie file. The torrent file contained metadata of the source file (movie file) and the URL of the BitTorrent tracker server.

To distribute the movie, the defendant sent the torrent file to several newsgroups. He then activated the torrent file on his computer, which caused his computer to connect to the tracker server. The tracker server queried the defendant's computer about the metadata of the torrent file. The tracker server then returned a list with the IP addresses of peer machines on the network and the percentages of the target file that existed on the peer machines.

Since the defendant's computer had a complete copy of the movie, the tracker server labeled it as a "seeder computer." The defendant maintained the connection between the tracker server and his computer so that other peers could download the movie from his computer.

4.2 Building the Model

The construction of a Bayesian network model begins with the main hypothesis that the analyst intends to determine. In order to prove the illegal act in the BitTorrent case, we use the following hypothesis:

H : The seized computer was used as the initial seeder to share the pirated file on a BitTorrent network.

Next, we express the possible states of the hypothesis (Yes, No and Uncertain) and assign probability values to these states. The values are also called the prior probabilities of the hypothesis.

Hypothesis H is the root node in the Bayesian network. Since it has no parent nodes, its prior probabilities are unconditional. To begin with, the probabilities of H are evenly distributed among its three states, i.e., $P(H) = (0.333, 0.333, 0.333)$ (Table 1).

Table 1. Prior probability of the root node.

Node	State	$P(H)$
H	Yes	0.333
	No	0.333
	Uncertain	0.333

Having established the root node, we proceed to explore evidence or events that are causally dependent on H . These are usually observable variables. However, note that sub-hypotheses may also be added under the root node. Although these sub-hypotheses do not have observable states, they are useful because they refine the model by producing a graph with more structure and increased clarity. Five sub-hypotheses are created to support the root hypothesis:

- H_1 : The pirated file was copied from the seized optical disk (found at the crime scene) to the seized computer.
- H_2 : A torrent file was created from the copied file.
- H_3 : The torrent file was sent to newsgroups for publishing.
- H_4 : The torrent file was activated, which caused the seized computer to connect to the tracker server.
- H_5 : The connection between the seized computer and the tracker server was maintained.

Table 2. Conditional probabilities of H_1 .

State	Yes	No	Uncertain
$H = \text{Yes}$	0.6	0.35	0.05
$H = \text{No}$	0.35	0.6	0.05
$H = \text{Uncertain}$	0.05	0.05	0.9

Since the sub-hypotheses are dependent on H , they are assigned conditional probability values. Table 2 presents the conditional probability values of Hypothesis H_1 given the state of H . Initial or prior probability

values are assigned to the possible states of H_1 for different states of H . For example, an initial value of 0.6 is assigned for the situation when H and H_1 are both Yes. This means that when the seized computer has been used as an initial seeder, the probability that the pirated file found on the computer had been copied from the optical disk seized at the crime scene is 0.6. However, it is also possible that, although the seized computer was the initial seeder, the pirated file was downloaded from the Internet or copied from another computer in a local network; a probability value of 0.35 is assigned to these scenarios.

Finally, there is the possibility that, even though the seized computer was the initial seeder, the evidence may not be able to confirm a Yes or No state for H_1 . Therefore, there is a chance that the seized computer was the initial seeder, but the source from where the pirated movie was copied is Uncertain.

Table 3. Conditional probabilities of H_2 , H_3 , H_4 and H_5 .

State	Yes	No	Uncertain
$H = \text{Yes}$	0.6	0.35	0.05
$H = \text{No}$	0.35	0.6	0.05
$H = \text{Uncertain}$	0.05	0.05	0.9

Table 3 presents the conditional probabilities of Hypotheses H_2 , H_3 , H_4 and H_5 given the state of H .

Following the assignment of conditional probabilities to the five sub-hypotheses, we proceed to develop the entailing casual events or evidence for the sub-hypotheses. This is because a Bayesian network propagates probabilities for linked hypotheses based on the states of events or evidence.

Hypothesis H and the five sub-hypotheses have a diverging connection. The nodes in a diverging connection influence each other when the state of their parent node is still unknown. Therefore, the five sub-hypotheses are related to each other in a probabilistic manner. Also, their probabilities are affected by all the child events or evidence under them.

To illustrate the Bayesian network methodology, we focus on Hypothesis H_1 : The pirated file was copied from the seized optical disk (found at the crime scene) to the seized computer.

5. Assigning Prior Probabilities

Items of digital evidence correspond to past digital events (or posterior evidence) that can be used to support or refute the five sub-hypotheses,




Figure 2. Partial Bayesian network for H_1 .

which, in turn, support or refute H . One of the main challenges in applying a Bayesian network to evaluate evidence is assigning probability values to posterior evidence. This is because the assignments are usually based on subjective personal beliefs. Although the personal beliefs (regarding a case) of a digital forensic analyst are assumed to arise from professional knowledge and experience, there is no means to determine whether they truly represent the accepted views of the digital forensic discipline, let alone whether or not the probability values assigned to posterior evidence are, in fact, accurate.

To enhance the reliability and accuracy of the probability assignments for posterior evidence, we attempted to use objective probability assignments obtained by aggregating the responses of experienced law enforcement agents and analysts. A questionnaire (available at www.cs.hku.hk/kylai/qr.pdf) was created to obtain the required information from personnel with the Technical Crime Bureau of the Hong Kong Police and the Computer Forensic Laboratory of Hong Kong Customs. The questionnaire solicited the following information from the respondents: (i) digital forensics training and experience, (ii) degree of belief in digital evidence resulting from general computer operations, and (iii) degree of belief in the digital evidence related to the operation of the BitTorrent protocol.

Responses were received from 31 law enforcement personnel. The weighted average approach was used to aggregate the probability values. For example, Item 7 of the questionnaire required respondents to gauge the probability range that the URLs and access times of web sites would be stored in the file named `index.dat` in the folder `History.IE5`. The answers received were: 20-40%: 1 respondent, 40-60%: 1 respondent, 60-80%: 6 respondents, 80-100%: 22 respondents, and Uncertain: 1 respondent. The weighted average of the probability of the Yes state was computed as: $(1 \times 0.3) + (1 \times 0.5) + (6 \times 0.7) + (22 \times 0.9) = 24.8$, which yielded a probability value $24.8/31 = 0.8$. The probability of the Uncertain state was computed as $1/31 = 0.03$. Therefore, the probability of the No state was $1 - 0.8 - 0.03 = 0.17$.

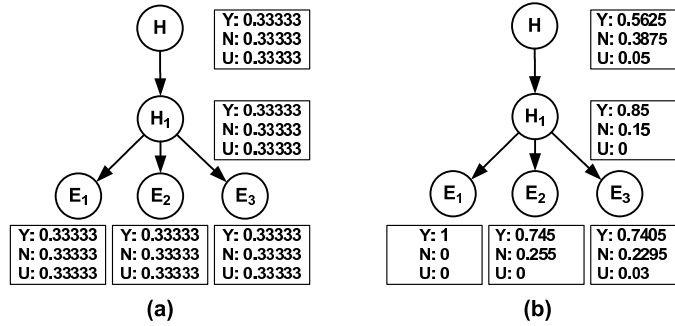


Figure 3. Probability values: (a) Initial; (b) Updated.

6. Analyzing Hypothesis H_1

The partial Bayesian network for Hypothesis H_1 is presented in Figure 2. The arguments describing events or evidence that would be caused by copying a file from an optical disk to a local hard disk are : (i) E_1 : Modification time of the destination file equals that of the source file (states: Yes, No, Uncertain), (ii) E_2 : Creation time of the destination file is after its own modification time (states: Yes, No, Uncertain) and (iii) E_3 : Hash value of the destination file matches that of the source file (states: Yes, No, Uncertain).

Table 4. Conditional probabilities of E_1 , E_2 and E_3

State	E_1			E_2			E_3		
	Y	N	U	Y	N	U	Y	N	U
$H = Y$	0.85	0.15	0	0.85	0.15	0	0.85	0.12	1.03
$H = N$	0.15	0.85	0	0.85	0.15	0	0.12	0.85	0.03
$H = U$	0	0	1	0	0	1	0.03	0.03	0.94

The next task is to assign conditional probability values to the events or evidence. Table 4 lists the conditional probabilities of E_1 , E_2 and E_3 , given the state of H_1 .

Next, the probability of H_1 based on the observed probabilities of E_1 , E_2 and E_3 is calculated. The MSBNx Bayesian Network Editor and Tool Kit [14] was used to calculate this probability and to propagate probability values within the Bayesian network.

The probability values for the network nodes are presented in Figure 3. Figure 3(a) presents the initial probability values in the network without any observed evidence. Figure 3(b) shows the updated probabilities

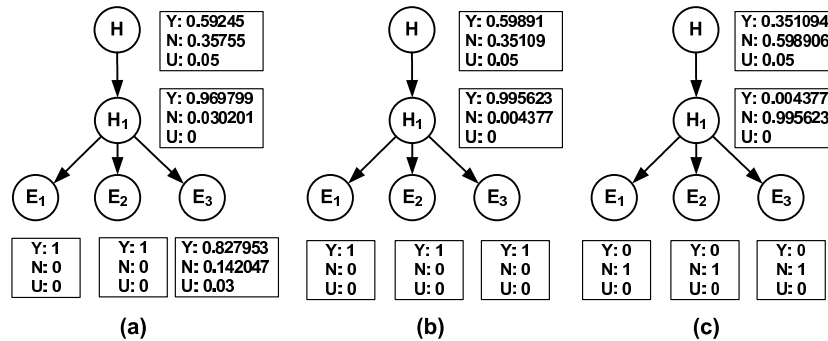


Figure 4. Propagated probability values.

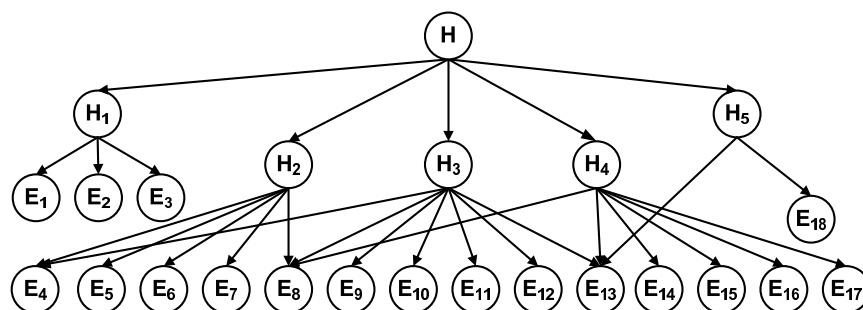
assuming that evidence E_1 is observed to be Yes while E_2 and E_3 are still unobservable.

Hypothesis H_1 is in a diverging connection with E_1 , E_2 and E_3 . Therefore, if the state of H_1 is unobserved, any change in the probability of E_1 will change the probability of H_1 . When H_1 changes, the likelihood ratios of E_2 and E_3 also change. Similarly, since H , H_1 and E_1 are in a serial connection, a change in the probability of E_1 will propagate to H if H_1 remains unobservable.

The Bayesian network has two more serial connections, $H \rightarrow H_1 \rightarrow E_2$ and $H \rightarrow H_1 \rightarrow E_3$. Therefore, any changes in the states of E_2 and E_3 will also affect the probabilities of H and H_1 .

Suppose we examine the state of the posterior evidence E_2 and find it to be Yes. The corresponding propagated probabilities in the network are shown in Figure 4(a). If the final posterior evidence E_3 is also observed to be Yes, the probabilities that result are shown in Figure 4(b). Note that when all the evidence states are Yes, the propagated probability for $H_1 = \text{Yes}$ is 99.6% and the corresponding probability for $H = \text{Yes}$ is 59.9%. In other words, if the states of E_1 , E_2 and E_3 are all Yes, then the digital forensic analyst can confirm that there is a 99.6% probability that H_1 (the pirated file was copied from the seized optical disk to the seized computer) is true. Furthermore, based on the 99.6% probability value for H_1 , the forensic analyst can conclude that H (the seized computer was used as the initial seeder to share the pirated file on a BitTorrent network) is true with probability 59.9%.

Figure 4(c) shows the resulting probabilities for the case where all the evidence states are No. The probability that H_1 is true drops to 0.4% and the probability that H_1 is false rises to 99.6%. Unless a posterior event or evidence exists, the probability that H is true drops to 35.1% and the probability that H is false rises to 59.9%.

**HYPOTHESES:**

- H** The seized computer was used as the initial seeder to share the pirated file on a BitTorrent network
- H₁** The pirated file was copied from the seized optical disk to the seized computer
- H₂** A torrent file was created from the copied file
- H₃** The torrent file was sent to newsgroups for publishing
- H₄** The torrent file was activated, which caused the seized computer to connect to the tracker server
- H₅** The connection between the seized computer and the tracker was maintained

EVIDENCE:

- E₁** Modification time of the destination file equals that of the source file
- E₂** Creation time of the destination file is after its own modification time
- E₃** Hash value of the destination file matches that of the source file
- E₄** BitTorrent client software is installed on the seized computer
- E₅** File link for the shared file is created
- E₆** Shared file exists on the hard disk
- E₇** Torrent file creation record is found
- E₈** Torrent file exists on the hard disk
- E₉** Peer connection information is found
- E₁₀** Tracker server login record is found
- E₁₁** Torrent file activation time is corroborated by its MAC time and link file
- E₁₂** Internet history record about publishing website is found
- E₁₃** Internet connection is available
- E₁₄** Cookie of the publishing website is found
- E₁₅** URL of the publishing website is stored in the web browser
- E₁₆** Web browser software is available
- E₁₇** Internet cache record about the publishing of the torrent file is found
- E₁₈** Internet history record about the tracker server connection is found

Figure 5. Bayesian network diagram.

7. Analyzing the BitTorrent Case

The overall Bayesian network diagram for the BitTorrent case is shown in Figure 5. When no observations are made on any entailing evidence, the initial probabilities of H_1 , H_2 , H_3 , H_4 and H_5 and, therefore, H are Yes = 33.33%, No = 33.33% and Uncertain = 33.33%.

Table 5. Probabilities of various hypotheses.

Hypothesis	(a)			(b)		
	Y(%)	N(%)	U(%)	Y(%)	N(%)	U(%)
H	92.54	7.45	0.01	92.27	7.72	0.01
H_1	99.71	0.29	0.00	99.70	0.30	0.00
H_2	99.98	0.0015	0.0185	99.92	0.07	0.01
H_3	99.98	0.02	0.00	99.80	2.20	0.00
H_4	99.93	0.07	0.00	99.51	0.49	0.00
H_5	89.31	10.47	0.22	99.45	10.33	0.22

When all the entailing evidence is switched to the state Yes, the propagated probabilities for the various hypotheses are as presented in Table 5(a).

Media reports about the BitTorrent trial mentioned that there was no indication that the torrent file was present on the seized computer. Also, there was no mention of cookies that are required to publish the torrent file in newsgroups. Therefore, the corresponding observations about the existence of the created torrent file (node E_8 in Figure 5) and cookies of newsgroups (node E_{14}) should be amended from Yes to No in order to reveal their impact on the hypotheses.

It is worth mentioning that the “torrent file node” (E_8) is a common node for H_2 , H_3 and H_4 . In other words, there is a converging connection to E_8 from these three hypotheses. According to the rules of probability propagation for a converging connection, when the state of E_8 is known, the probabilities of H_2 , H_3 and H_4 will influence each other. Therefore, a change in the state of E_8 changes the probabilities of these three hypotheses.

Furthermore, since H_1 , H_2 , H_3 , H_4 and H_5 are in a diverging connection with the parent hypothesis H , changes to the probabilities of H_2 , H_3 and H_4 influence the probabilities of H_1 and H_5 . Table 5(b) shows the probability values obtained after the states of E_8 and E_{14} are changed from Yes to No. The propagated probability for H from the available evidence is 92.27%. In other words, based on the observed evidence, there is a probability of 92.27% that the seized computer was used as the initial seeder to distribute the pirated movie on a BitTorrent

network. This is the most that a digital forensic analyst can provide. It is up to the court to decide whether or not this probability value is sufficient to support the Hypothesis H .

Note that other evidence exists in the BitTorrent case. This includes email exchanges, detailed comparisons of the torrent file metadata with computer trails, and timeline analysis. However, as the focus of this paper is to demonstrate the utility of Bayesian networks in digital forensic investigations, only the most important pieces of digital evidence were considered in the discussion.

8. Conclusions

A Bayesian network is a useful formalism for quantifying and propagating the strengths of investigative hypotheses and supporting evidence. The Internet piracy trial provides an excellent case study for validating the approach. The hypotheses in the case and their supporting events and evidence are clearly specified, along with their causal relationships and probability values. Thus, the Bayesian network model is not only an analytical tool for evaluating evidence, but also a tracking tool that enables digital forensic practitioners to review and analyze the original findings.

The subjectivity involved in assigning probabilities can be alleviated to some extent by using a survey instrument and aggregating the responses obtained from expert investigators. However, it is difficult to completely eliminate the subjective aspects, especially with regard to the assignment of prior probabilities to posterior evidence. Our future research will investigate this aspect in more detail with the goal of enhancing the accuracy, precision and reliability of the Bayesian network model for reasoning about digital evidence.

References

- [1] C. Aitken and F. Taroni, *Statistics and the Evaluation of Evidence for Forensic Scientists*, John Wiley and Sons, New York, 2004.
- [2] V. Baryamureeba and F. Tushabe, The enhanced digital investigation process model, *Proceedings of the Fourth Digital Forensic Research Workshop*, 2004.
- [3] S. Ciardhuain, An extended model of cybercrime investigations, *International Journal of Digital Evidence*, vol. 3(1), 2004.
- [4] R. Cook, I. Evett, G. Jackson, P. Jones and J. Lambert, A model for case assessment and interpretation, *Science and Justice*, vol. 38, pp. 151–156, 1998.

- [5] R. Cowell, Introduction to inference for Bayesian networks, *Proceedings of the NATO Advanced Study Institute on Learning in Graphical Models*, pp. 9–26, 1998.
- [6] P. Good, *Applying Statistics in the Courtroom: A New Approach for Attorneys and Expert Witnesses*, Chapman and Hall/CRC Press, Boca Raton, Florida, 2001.
- [7] International Association of Computer Investigative Specialists, Forensic procedures, Fairmont, West Virginia (www.cops.org/for_ensicprocedures), 2007.
- [8] International Organization on Computer Evidence, International principles for computer evidence, *Forensic Science Communications*, vol. 2(2), 2000.
- [9] J. Jones, Y. Xiang and S. Joseph, Bayesian probabilistic reasoning in design, *Proceedings of the IEEE Pacific Rim Conference on Communications, Computers and Signal Processing*, pp. 501–504, 1993.
- [10] K. Kent, S. Chevalier, T. Grance and H. Dang, Guide to Integrating Forensic Techniques into Incident Response, Special Publication 800-86, National Institute of Standards and Technology, Gaithersburg, Maryland, 2006.
- [11] J. Keppens and J. Zeleznikow, A model based reasoning approach for generating plausible crime scenarios from evidence, *Proceedings of the Ninth International Conference on Artificial Intelligence and Law*, pp. 51–59, 2003.
- [12] R. Loui, J. Norman, J. Altepeter, D. Pinkard, D. Craven, J. Lindsay and M. Foltz, Progress on Room 5: A testbed for public interactive semi-formal legal argumentation, *Proceedings of the Sixth International Conference on Artificial Intelligence and Law*, pp. 207–214, 1997.
- [13] Magistrates' Court at Tuen Mun, Hong Kong Special Administrative Region v. Chan Nai Ming, TMCC 1268/2005, Hong Kong, China (www.hklii.hk/hk/jud/en/hksc/2005/TMCC001268A_2005.html), 2005.
- [14] Microsoft Research, MSBNx: Bayesian Network Editor and Tool Kit, Microsoft Corporation, Redmond, Washington (research.microsoft.com/adapt/MSBNx).
- [15] J. Mortera, A. Dawid and S. Lauritzen, Probabilistic expert systems for DNA mixture profiling, *Theoretical Population Biology*, vol. 63(3), pp. 191–206, 2003.

- [16] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*, Morgan Kaufmann, San Mateo, California, 1988.
- [17] S. Peisert, M. Bishop, S. Karin, M. Bishop and K. Marzullo, Principles-driven forensic analysis, *Proceedings of the New Security Paradigms Workshop*, pp. 85–93, 2005.
- [18] D. Poole, Probabilistic Horn abduction and Bayesian networks, *Artificial Intelligence*, vol. 64(1), pp. 81–129, 1993.
- [19] H. Prakken, C. Reed and D. Walton, Argumentation schemes and generalizations in reasoning about evidence, *Proceedings of the Ninth International Conference on Artificial Intelligence and Law*, pp. 32–41, 2003.
- [20] D. Walton, Argumentation and theory of evidence, in *New Trends in Criminal Investigation and Evidence – Volume II*, C. Breur, M. Kommer, J. Nijboer and J. Reijntjes (Eds.), Intersentia, Antwerp, Belgium, pp. 711–732, 2000.