AI&Law manuscript No.
(will be inserted by the editor)

Building Bayesian Networks for Legal Evidence with Narratives

A case study evaluation

**Charlotte S. Vlek · Henry Prakken ·
Silja Renooij · Bart Verheij**

the date of receipt and acceptance should be inserted later

Abstract In a criminal trial, evidence is used to draw conclusions about what happened concerning a supposed crime. Traditionally, the three main approaches to modeling reasoning with evidence are argumentative, narrative and probabilistic approaches. Integrating these three approaches could arguably enhance the communication between an expert and a judge or jury. In previous work, techniques were proposed to represent narratives in a Bayesian network and to use narratives as a basis for systematizing the construction of a Bayesian network for a legal case. In this paper, these techniques are combined to form a design method for constructing a Bayesian network based on narratives. This design method is evaluated by means of an extensive case study concerning the notorious Dutch case of the Anjum murders.

Keywords Legal reasoning · Bayesian networks · Narrative

C.S. Vlek
Institute of Artificial Intelligence, University of Groningen
Tel.: +31 50 363 6399, E-mail: c.s.vlek@rug.nl

H. Prakken
Department of Information and Computing Sciences, Utrecht University
And Faculty of Law, University of Groningen
E-mail: h.prakken@uu.nl

S. Renooij
Department of Information and Computing Sciences, Utrecht University
E-mail: s.renooij@uu.nl

B. Verheij
Institute of Artificial Intelligence, University of Groningen
CodeX , Stanford University
E-mail: b.verheij@ai.rug.nl

1 Introduction

In a criminal trial, the available evidence is used to draw conclusions about the events that took place (referred to in legal terminology as the *facts* of the crime). In the formal study of rational legal reasoning about evidence, drawing conclusions from evidence is usually formalized by one of three approaches: argumentative, narrative or probabilistic reasoning (Kaptein et al, 2009). While a judge or jury might be more inclined to use argumentation or narratives, a forensic expert might use probabilities to report his findings. In this paper a method is presented to integrate narrative and probabilistic approaches.

Our method aims to enhance the communication between an expert and a judge or jury by modeling the evidence and a number of relevant scenarios in a Bayesian network. Bayesian networks have been studied as a tool for probabilistic reasoning about legal cases (Hepler et al, 2004; Keppens, 2011; Lagnado et al, 2013; Taroni et al, 2006), although the use of Bayesian methods in court remains subject of debate (see Fenton and Neil (2012) but also Berger and Aben (2010b,c,a)). Bayesian networks provide a solid mathematical framework that can capture the quantitative information supplied by forensic experts, while the graphical part can model qualitative information about the connections between variables in the case as a whole. This forms a first step towards a better communication: while the Bayesian network can be used to store both qualitative and quantitative information, methods for communicating this in an informative way to a judge or jury still need to be developed. We hope to address this in future work.

In previous papers (Vlek et al, 2013a,b,c), we discussed techniques for representing narratives in a Bayesian network and for building Bayesian networks based on narratives. In this paper, these techniques are combined to form one design method for building Bayesian networks for reasoning with evidence in the legal domain. A main goal of this paper is to evaluate the design method by means of a case study. The aim of our design method is to combine the best of two worlds: while Bayesian networks enable a formalization of both qualitative (the graphical structure) and quantitative (the numbers) aspects of the case, narratives provide the context needed for finding all relevant variables for the network and for building the structure.

In the narrative approach to reasoning with evidence, a crime scenario forms a context for the evidence. A scenario or story is a collection of states and events with some coherent structure (Bex, 2011). When investigating a crime, scenarios are formed to make sense of the evidence (Poot et al, 2004), and as Pennington and Hastie (1992) showed, jurors use stories or scenarios to organize the evidential data. Wagenaar, Van Koppen and Crombag (1993) claim that in any legal case, various scenarios should be considered in order to avoid tunnel vision.

Among forensic scientists, Bayesian networks have become popular as a general method for reasoning with evidence (Hepler et al, 2004; Keppens, 2011; Lagnado et al, 2013; Taroni et al, 2006). A Bayesian network is a representation of a joint probability distribution. A graph represents the variables and their (in)dependencies, while probability tables specify conditional probabilities for each variable. From the network, any probability of interest can be calculated, including updated probabilities after evidence has been added. In legal applications, a Bayesian network is commonly used to model the (in)dependencies between a hypothesis and the evidence. Supplying the appropriate numbers for the probability tables is a known limitation in applications of Bayesian networks, but a number of elicitation techniques exist (Renooij, 2001). The focus of this paper is not on the numbers, but on methods for building the graphical structure of a Bayesian network.

In the legal domain, every case is different, demanding a custom-made Bayesian network. A number of researchers have worked on the development of techniques that can help systematize the process of building a Bayesian network. For legal applications, Hepler, Dawid and Leucari (2004) proposed to look at recurrent substructures that can be used as building blocks throughout various networks. Fenton, Neil and Lagnado (2013) call such substructures *idioms*, and propose a list of legal idioms that can be used when constructing a network for a case. These idioms are helpful in finding the local structure for a network, including the connections between specific types of variables. However, a method for determining which variables are relevant to a case and therefore which variables to include in a network is lacking.

In this paper, we propose to use narratives as a basis for building a Bayesian network. This results in a method that takes scenarios as a starting point to decide which variables are relevant to a case. Furthermore, representing narratives in a Bayesian network takes us one step closer to an integration of the various approaches to reasoning with evidence. The methods proposed in this paper currently assume that the police investigation produced a number of relevant scenarios. The problem that is addressed is thus how these existing scenarios can be represented and evaluated with probabilistic techniques.

To enable the representation of a crime scenario in a network, in (Vlek et al, 2013a,b,c) we extended the list of idioms by Fenton et al. with four narrative idioms. For building a Bayesian network, in (Vlek et al, 2013c) we introduced the method of unfolding a scenario, using narratives as a basis to gradually construct a graph. In this paper, these techniques are taken together to form a design method for building Bayesian networks for legal cases.

The design method aims to alleviate three common difficulties in reasoning with evidence: (1) tunnel vision, (2) the problem of a good story pushing out a true story and (3) finding the relevant variables for a model of the case. An advantage of the narrative approach is the use of multiple stories or scenarios to prevent tunnel vision, as described above. By explicitly taking multiple scenarios into account, the narrative approach in our design method aims to prevent (1). A disadvantage of the narrative approach is the problem of a good story being chosen over a true story (Pennington and Hastie, 1993). By integrating a narrative and probabilistic approach, the various scenarios in a case can be compared in their formal representation, the Bayesian network, such that the more likely story can be chosen over what may sound as a good story, addressing (2). Finally, the existing probabilistic approaches to modeling legal cases often do not describe how to find which variables are relevant. By taking narratives as a starting point, heuristics for finding the variables can be formulated, addressing (3).

A main goal of this paper is to evaluate our design method by means of a case study, to show how our approach handles a complex case. To this end, an analysis, using our design method, of the so-called Anjum case (Crombag and Israëls, 2008) and (Bex, 2011, Chapter 6) is included as an evaluation. This meets Conrad and Zeleznikow's (2013) warning that without an evaluation, "no researcher can expect the broader audience to be convinced of the benefits and utility of their work". The results of the case study will show that our design method is indeed capable of representing narratives in a Bayesian network. The method thereby includes features of the narrative approach that help to prevent tunnel vision, while formalizing the narrative approach to the effect that a good story will not push out a true story. Finally, the use of narratives as a basis turns out to be helpful when gradually constructing the Bayesian network for a complex case.

Our choice to formalize narratives with the use of Bayesian networks comes with the drawback of having to specify all the required probabilities. This well-known limitation will be encountered in our case study as well, where the design method requires many numbers to be made precise. The elicitation of probabilities for the case study will be discussed in Section 4.6. In the narrative idioms in Section 3.1, certain conditional probabilities will be fixed since they follow from the existing logical relations, such as that of an event being an element of a scenario. Other probabilities will have to be determined based on a user's knowledge of the world. In our design method, as expected from a Bayesian approach, many probabilities in the network will be a subjective representation of the real world. Nonetheless, by making these numbers precise, the subjective interpretation of a legal case is made explicit. Therefore, our design method is intended to formalize a subjective account of a legal case, thus forming a tool for a judge or jury to structure their thoughts rather than a tool for reaching an objective verdict.

The contributions of this paper are twofold: (1) continuing from our previous work, a design method for building a Bayesian network from scenarios is presented and (2) the method is evaluated by means of an extensive case study. The remainder of this paper is organized as follows. First, some preliminaries are discussed in Section 2. Then the design method including four narrative idioms and the concept of unfolding is presented in Section 3. In Section 4 the case study is presented, followed by a discussion in Section 5. The paper ends with a discussion of related work in Section 6 and a conclusion in Section 7.

2 Preliminaries

A Bayesian network is a compact representation of a joint probability distribution (JPD) (Jensen and Nielsen, 2007). It consists of a directed acyclic graph (DAG) with a set of nodes representing variables in the domain and a set of

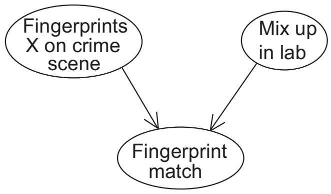

(a) Example of a Bayesian network structure

FP X = true	FP X = false
0.2	0.8

(b) The prior probability P(`Fingerprints X on crime scene`)

Mix up = true	Mix up = false
0.001	0.999

(c) The prior probability P(`Mix up in lab`)

	Mix up = true		Mix up = false	
	FP X = true	FP X = false	FP X = true	FP X = false
match = true	0.01	0.01	0.999	0.001
match = false	0.99	0.99	0.001	0.999

(d) CPT for P(`Fingerprint match`|`Fingerprints X on crime scene, Mix up in lab`)

Fig. 1: An example of a Bayesian network: graph and probability tables

directed edges between the nodes representing (possible) correlation between the variables. The graph as a whole describes the probabilistic independence relations on the modeled variables. The variables each have a finite set of mutually exclusive states. Each node has a conditional probability table (CPT) describing for a node A the probability distribution conditioned on combinations of values for its parents B_1 ... B_n in the graph (nodes with an arrow pointing to A), that is $P(A|B_1, ..., B_n)$.

A simple example of a Bayesian network is shown in Figure 1. The graph has three nodes, representing that a person X left fingerprints at the crime scene (`Fingerprints X on crime scene`), a fingerprint match was found with fingerprints from the crime scene and person X (`Fingerprint match`) and there was a mix up in the lab (`Mix up in lab`). In this example the nodes have binary values: each of these nodes has values true and false. The probability tables in Figure 1 specify the underlying probabilities: from Table 1d it can be read that the probability that a match would be found with person X, when X left no fingerprints at the crime scene and there was no mix up, is very small: 0.001.

Whenever the value of a node is observed (evidence is found), this can be entered in the network: we say that a node is *instantiated* to the corresponding value. Inference in the network leads to posterior probabilities given these values. Instantiating nodes may lead to changed (in)dependencies between variables, depending on the type of connections. Independences that hold in the distribution captured by the Bayesian network can to some extent be read from the network's graph, using the notion of *d-separation*. When two variables A and B are d-separated, a change in the certainty of A has no influence on the certainty of B and vice versa.

Whether two variables A and B are d-separated, depends on how they are connected to each other via a chain of nodes and arrows. Crucial is whether this chain includes a head-to-head node (a node with two incoming arrows in the chain). Such a *converging* connection is shown in the example in Figure 1a. The upper two nodes are connected via a head-to-head connection in `Fingerprint match`.

For a converging connection with arrows meeting head-to-head in V, the chain is *inactive* as long as neither V, nor any of V's descendants, are instantiated. When there are no active chains between A and B, they are d-separated. As soon as V or any of its descendants is instantiated, the chain between A and B becomes active, and A and B are said to be d-connected. In the example in Figure 1a this means that person X leaving fingerprints at the crime scene has no influence on a possible mix up in the lab as long as no information is available about a possible match. However, once the match is found, there is an active path and the variables become d-connected: knowing that the match is a result of a mix up will make the probability of the person X actually leaving fingerprints at the crime scene drop.

When a chain of arrows between A and B has no head-to-head connections, the connection is serial. The chain is now active whenever *none* of the nodes are instantiated (which differs from the converging case). The nodes A and B are thus d-connected when none of the nodes in the chain have been instantiated. When a node V in the chain is instantiated, the chain becomes inactive. When there are no active chains connecting A and B, they are d-separated.

The formal definition of d-separation is as follows:

Definition 1 Two variables A and B are d-separated given a set of observed nodes E when *for all* chains between A and B there is an intermediate variable V (not equal to A or B) such that either

- the connection from A to B via V is serial and $V \in E$, or
- the connection from A to B via V is converging and $V \notin E$, nor are any of the descendants of V in E.

When two variables are not d-separated, they are d-connected.

The concept of d-separation allows for reading independences from a Bayesian network's directed graph. If variables A and B are d-separated given V, then A and B are independent given V in the JPD represented by the network. The JPD $P(U)$, given by

$$P(U) = \Pi_{i=1}^{n} P(A_i|\mathrm{pa}(A_i))$$

where A_1 to A_n are the nodes in the network and $\mathrm{pa}(A_i)$ are the parents of A_i, respects these independences.

Various tools are available for working with Bayesian networks, such as GeNIe 2.0.[1] Such tools can be used to calculate any prior or posterior probability of interest from the network, such as for example, the probability that person X left fingerprints at the crime scene given that a fingerprint match was found in Figure 1.

[1] GeNIe 2.0 is available for free on genie.sis.pitt.edu

3 A design method for forensic Bayesian networks

In this section our design method for building Bayesian networks for legal cases is presented. It employs a number of legal idioms, including idioms from Fenton et al (2013), Lagnado et al (2013) and four newly developed narrative idioms (which will be presented in Section 3.1). The building process relies on the concept of unfolding a scenario, which will be introduced in Section 3.2. Finally, in Section 3.3 a protocol is presented, describing which steps to take when constructing a Bayesian network with our design method.

3.1 Narrative idioms

An idiom is a recurring structure that can be used as a building block in various Bayesian networks for different cases. For the particular application of representing narratives in a Bayesian network, we developed four narrative idioms in (Vlek et al, 2013a,b,c).

The aim of the narrative idioms in particular is to capture the notion of coherence of a scenario. Consider the following scenario for a burglary: A intended to steal something. Therefore, A broke into a house and took some items. Note that this scenario is more than just a collection of events: the elements together form a coherent whole. To see why the burglary scenario is coherent, suppose we come home to find a broken window or a forced door. Immediately, we imagine that a burglary must have taken place and we assume that some items must be missing as well.

In the narrative field, coherence is considered a key property of a story or scenario (e.g. Pennington and Hastie (1993)). However, finding a clear definition of coherence in the scientific literature on narrative proves to be difficult. Various researchers have studied the idea of stories following some pattern, such as story grammars (Rumelhart, 1975), scripts (Schank and Abelson, 1977) or schemes (Pennington and Hastie, 1993). By following such a pattern, a story can be assured to 'have all of its parts', thereby forming a coherent whole.

In this text we are not concerned with the fundamental question of why a scenario is coherent, but rather with how we can capture the existing coherence in our models. In particular, we intend to model the phenomenon that was described in the burglary example above, where evidence for part of the scenario affects the scenario as a whole. This is called *transfer of evidential support* (see also Bex (2011)). By supplying evidence for one element of a scenario, not only does this influence the probability of that particular element of the scenario, but it also affects the probability of the entire scenario. The scenario idiom and the subscenario idiom are specifically designed such that they capture the transfer of evidential support within a scenario.

The idioms in the following subsections are meant to model a scenario as a whole (the scenario idiom), a subscenario (the subscenario idiom) and small variations within the scenario (the variation idiom). Finally, the merged scenarios idiom is used to obtain one network modeling all scenarios.

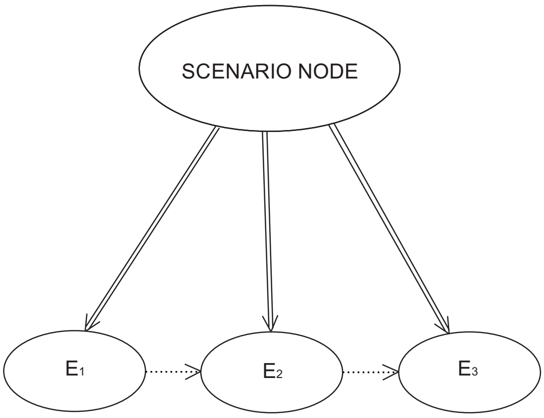

Fig. 2: The scenario idiom. Dotted lines suggest possible connections between elements of the scenario, including multiple dependencies (not shown). Double arrows represent connections where the probabilities are partially fixed by logical relations.

Table 1: Probability table for an element of the scenario

	ScN = true	ScN = false
E_i = true	1	...
E_i = false	0	...

In this section, these four narrative idioms are discussed and illustrated with examples from a burglary case.

3.1.1 Scenario idiom

With the scenario idiom, a scenario can be captured in a Bayesian network. The main feature of the scenario idiom is that it captures the coherence of a scenario. This is done by connecting all elements of a scenario by means of a *scenario node* with outgoing arcs only, pointing from the scenario node to each element of that scenario. Via this node, as long as it remains uninstantiated, a transfer of evidential support is guaranteed.

A general version of the scenario idiom is shown in Figure 2. An instance of the scenario idiom consists of the following:

1. A set of n element nodes $E = \{E_1, \ldots..., E_n\}$ with values true and false for each element of the scenario,
2. arrows between the element nodes whenever there is a connection between the corresponding elements of the scenario (as dictated by the formalization of a Bayesian network), and

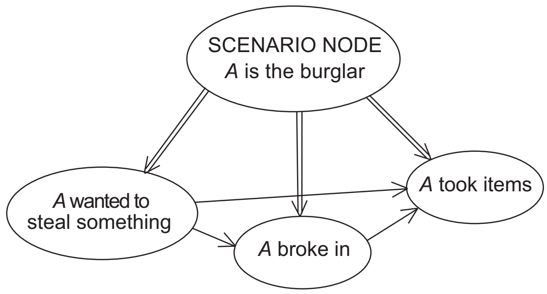

Fig. 3: A scenario for the burglary case

3. a scenario node, with arrows from the scenario node to each element node $E_i \in E$.

Note that the resulting graph will always be acyclic (assuming that the connections between elements can be modeled acyclic) because the scenario node only has outgoing nodes.

Modeling the burglary scenario described in the introduction above with our scenario idiom amounts to a directed acyclic graph as shown in Figure 3. Since there is an arrow pointing from the scenario node to any element of the scenario, there is an active path via the scenario node for any pair of element nodes (note that the value of the scenario node is never observed and therefore never instantiated) so they are always d-connected. This allows the probability of all elements of the scenario to change as one element changes in probability, modeling the desired transfer of evidential support.

The underlying probabilities for the scenario idiom are determined according to two principles:

1. whenever the scenario is true, all of its elements must be true and
2. when the scenario is not true, elements of it might still be true.

Because of (1), several but not all numbers in the probability tables for a scenario are fixed by the existing logical relations. Conditional probabilities for connections between elements in the scenario need to be determined separately, as they depend on the scenario. Eliciting these probabilities is certainly not an easy task, see Renooij (2001) for elicitation methods. Any element node $E_i \in E$ has as parents $\mathrm{pa}_E(E_i) = \mathrm{pa}(E_i) \cap E$ (parents among other element nodes) and the scenario node. For such a node, we fix the following probabilities for any assignment of $\mathrm{pa}_E(E_i)$:

– $P(E_i \texttt{ = true} | \texttt{Scenario node= true}, \mathrm{pa}_E(E_i)) = 1$
– $P(E_i \texttt{ = false} | \texttt{Scenario node= true}, \mathrm{pa}_E(E_i)) = 0$

This is shown in Table 1 for an element node with only the scenario node as a parent. The conditional probabilities when the scenario node does not hold depend on a user's subjective account of the case (shown as dots in Table 1). In our figures, double arrows indicate a partly deterministic relation. These

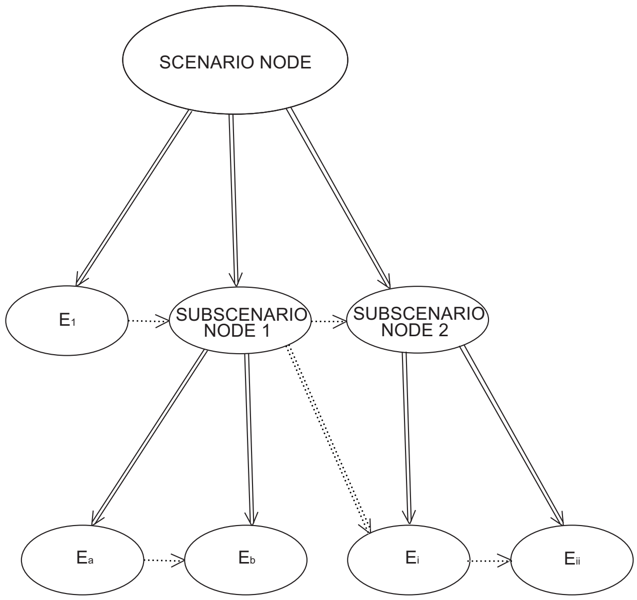

Fig. 4: The subscenario idiom. Dotted lines suggest possible connections between elements of the scenario and double arrows represent connections where the probabilities are partially fixed by logical relations. The double dotted line suggests that an element can be part of more than one scenario.

arrows are just meant to be a visual aid; there is no technical difference between the double and single arrows.

Finally, the plausibility of a scenario is specified in the probability table for the scenario node. Plausibility, a term from narrative research, is one of the key factors that determine the quality of a scenario, according to Pennington and Hastie (1993). When a story or scenario is found to be credible without taking any evidence into consideration, it is said to be plausible. This notion therefore arguably translates to the prior probability of the scenario as a whole, which is represented in the probability table for the scenario node.

3.1.2 Subscenario idiom

In many cases, a scenario consists of subscenarios: smaller stories that together form a part of the larger story. For example, in the scenario from Figure 3, the event '*A* broke in' can be replaced with a subscenario describing just how this happened: *A* broke the window and went into the house. For modeling such subscenarios, the subscenario idiom was developed.

The subscenario idiom has the same general structure as the scenario idiom, but is always modeled as part of a larger scenario idiom. Again the elements of the subscenario are represented with binary nodes and a subscenario node

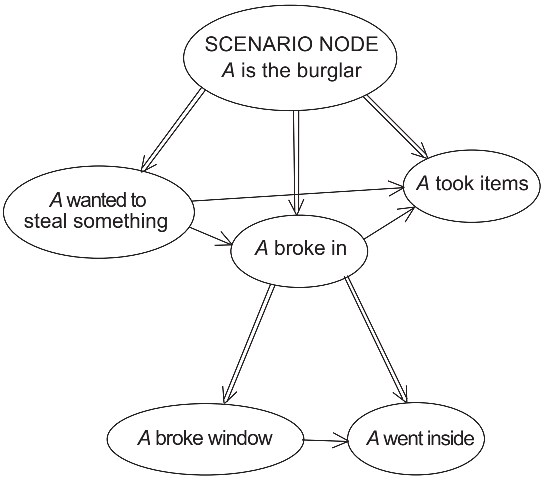

Fig. 5: A subscenario for the burglary case

is included at the top with arrows pointing to all elements of the subscenario, resulting in a directed acyclic graph. Probabilities are determined according to similar principles: whenever the subscenario node is true, it follows that the elements of that subscenario must be true, but not the other way around. A subscenario is always part of a scenario, as shown in the general version of the subscenario idiom in Figure 4 and in Figure 5, modeling the example from the burglary case. Therefore, the probability table for the subscenario node does not contain prior probabilities, like the scenario node. Instead, the CPT for the subscenario node expresses the probabilities for the event it describes within the scenario, including the fixed probabilities that follow from the logical relations in the scenario idiom. As a result, by replacing the node 'A broke in' with a subscenario, the CPT of the subscenario node is equal to the CPT for 'A broke in'.

3.1.3 Variation idiom

The subscenario from Section 3.1.2 (see Figure 5 for this idiom) described how A broke into the house. However, suppose that it is clear that this person broke in, but there are three possibilities: either A smashed a window with a stone, A forced a door, or A picked a lock to get in. By working out these three possibilities, evidence about each possibility can be included in the network. Rather than modeling this with three separate instances of the scenario idiom, the variations can best be modeled within one scenario. Reasons to do the latter are (1) that it will reduce the complexity of the graph (since no new scenario idioms are needed) and (2) that it will help the modeler to maintain overview of various scenarios. The variation idiom can model small variations within a scenario (see Figure 6). The idiom is particularly suited for variations

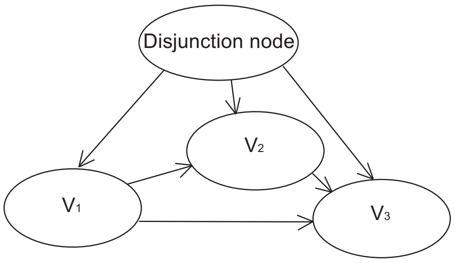

Fig. 6: The variation idiom with three variations. Connections between variations are needed to make sure that exactly one of the variations holds whenever the disjunction node is true.

that do not influence the overall conclusion of the scenario (i.e. that A broke in).

Typically, a variation would be modeled in a Bayesian network as multiple values of one node (for example, a node ‘forced entry’ with values ‘smashed window’, ‘forced door’ and ‘picked lock’). However, in our variation idiom, the variations are modeled as binary nodes (with values true/false). This was done to enable a situation where each variation is itself a subscenario. This has the advantage that not only multiple variations of forced entry (smashed window/forced door/picked lock) can be included, but entire subscenarios can serve as variations, such as the subscenario describing how A smashed the window with a stone, resulting in the window breaking and another subscenario describing how A forced the door.

The variation idiom (see Figure 6 for an instance with three variations) consists of a set of n variation nodes $V = \{V_1 \ ... \ V_n\}$ for all n variations together with a disjunction node. Arrows point from the disjunction node to each variation V_i and for any pair of variations V_i and V_j there is an arrow $V_i \rightarrow V_j$ if and only if $i < j$. Since arrows point only from the disjunction node to a variation and from a variation with a lower index to a variation with a higher index, the result will be acyclic. Connecting arrows between any pair of variations are needed to be able to express that when the disjunction node is true, at least one variation V_i must be true.

For the burglary example, the disjunction node can be written as A `broke window/forced door/picked lock`. There will be three variation nodes A `broke window` and A `forced door` and A `picked lock`. Arrows point from the disjunction node to each variation, and for the connecting arrows between variations there are arrows A `broke window` to A `forced door`, from A `broke window` to A `picked lock` and from A `forced door` to A `picked lock`. In Figure 7, a scenario is shown with the window/door/lock variation modeled with the variation idiom. Note that the window variation is a small subscenario. The disjunction node is also an element node in the larger scenario: when the scenario node is true, it follows that the element ‘A broke

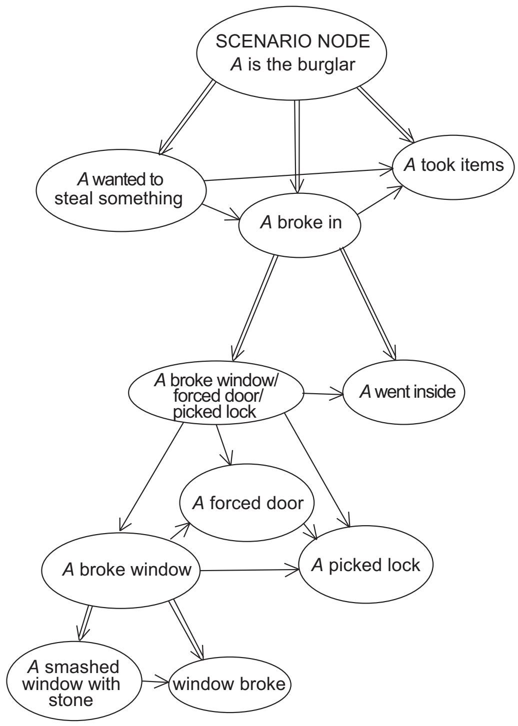

Fig. 7: A scenario with a variation

window/forced door/picked lock' must be true, without specifying which of the three variations is the case.

Note that the arrows from the disjunction node to each variation node V_i are directed from disjunction to variation. Although arrows from variation to disjunction might result in a simpler structure (in that case, no connecting arrow is needed between the variations), it is impossible to model variations within a scenario idiom this way. This is because the disjunction node would then have incoming arrows from the scenario node and all of the variations, resulting in the following probability that needs to be specified in the conditional probability table of the disjunction node:

$$P(\texttt{Disjunction node = true}|\texttt{Scenario Node = true}, V_1, ..., V_n \texttt{ = false})$$

Since we want the disjunction node to behave like an element node in the scenario idiom, the probabilities for this node need to be as shown in Table 1. This means that the above probability needs to be set to 1: when the scenario is true, it follows that all elements must be true. However, in order to model that the disjunction node is false when all variations are false, the above probability would need to be 0. Therefore, a structure with arrows from

Table 2: Conditional probability table for variation 1.

	`disjunction node = true`	`disjunction node = false`
V_1 `= true`	0.33	0
V_1 `= false`	0.67	1

Table 3: Conditional probability table for variation 2.

	`disjunction node = true`		`disjunction node = false`	
	V_1 `= true`	V_1 `= false`	V_1 `= true`	V_1 `= false`
V_2 `= true`	0	0.5	0	0
V_2 `= false`	1	0.5	1	1

Table 4: Conditional probability table for variation 3.

	`disjunction node = true`				`disjunction node = false`			
	V_1 `= true`		V_1 `= false`		V_1 `= true`		V_1 `= false`	
	V_2`=T`	V_2`=F`	V_2`=T`	V_2`=F`	V_2`=T`	V_2`=F`	V_2`=T`	V_2`=F`
V_3 `= T`	0	0	0	1	0	0	0	0
V_3 `= F`	1	1	1	0	1	1	1	1

variations to disjunction node cannot capture the variations as desired while leaving the scenario idiom intact.

With the variation idiom as shown in Figure 6, the scenario of which the variation is a part remains a coherent whole: the scenario idiom remains intact. With this structure the connecting arrows between variations are needed to model that exactly one of them must be true. This is done by specifying the probability tables as follows:

- If the disjunction node is false, the disjunction as a whole is false so none of the variations can be true. This leads to the following probability for any V_i and any assignment of $V_1, \ldots, V_{i-1}$: $P(V_i\ \texttt{= true}|\texttt{disjunction node = false}, V_1, \ldots, V_{i-1}) = 0$;
- If the disjunction node is true, this means that exactly one of the variations must hold.
  - To make sure that at least one variation holds, the last variation V_n (out of n variations) will be true with probability 1 when all other variations are false. Note that for V_n, all other variations V_i with $i < n$ are parents. Therefore, in the probability table it can be specified that $P(V_n\ \texttt{= true}|\ \texttt{disjunction node = true}, V_1\texttt{= false} \ldots V_{n-1}\ \texttt{= false}) = 1$.
  - To make sure that at most one variation holds, for a variation V_i, when there is some V_j`= true` with $j < i$, then V_i must be false. Therefore, $P(V_i\ \texttt{= true}|\ \texttt{disjunction node = true}, V_1\texttt{= false} \ldots V_j\texttt{= true} \ldots V_{i-1}\texttt{= false}) = 0$.

Table 5: CPT for the constraint node on scenario node 1 (`ScN1`) and scenario node 2 (`ScN2`)

	ScN1 = true		ScN1 = false	
	ScN2 = true	ScN2 = false	ScN2 = true	ScN2 = false
Constraint = scn 1	0	1	0	0
Constraint = scn 2	0	0	1	0
Constraint = NA	1	0	0	1

– When there is no further information available about the probabilities of the variations, they can be modeled to be equally likely.[2] Then, the probability of a variation V_i to be true, given that all variations with a lower index V_j with $j < i$ are false, is set to $1/(n-i+1)$. So $P(V_i\ \texttt{= true}|\ \texttt{disjunction node = true}, V_1\texttt{= false} \ldots V_{i-1}\texttt{= false})$ $= \frac{1}{n-i+1}$.

An example of these probabilities for the situation with three variations is shown in Tables 2, 3 and 4. Note that the probability table also contains numbers for $P(V_2\ = \texttt{true}|\texttt{disjunction node = false}, V_1\ = \texttt{true})$, expressing the probability that variation 2 occurs given that the disjunction is false but variation 1 is true, which should be impossible since the disjunction being false would yield variation 1 to be false as well. This situation will indeed never occur, since the probability table for variation 1 guarantees that variation 1 is false whenever the disjunction node does not hold. Hence, these numbers in the probability table for variation 2 are really undefined, but can be given arbitrary values to enable calculations in the Bayesian network.

The disjunction node is an element of the scenario, so the probability table for the disjunction node expresses the probability of this element, representing either of the variations.

3.1.4 Merged scenarios idiom

By representing multiple scenarios in one Bayesian network, various scenarios can be compared. In order to model all scenarios in one network, the merged scenarios idiom can be used. This idiom puts a constraint on the scenario nodes of separate scenarios, making sure that exactly one scenario can be true. By modeling small variations with the variation idiom, the scenarios that are modeled separately are really different.

Note that the merged scenarios idiom assumes that all scenarios in it are mutually exclusive. This means that the merged scenarios idiom should only be applied to scenarios for which this is the case. As an example, consider a burglary case where fingerprints of A were found, and footsteps of B. Three

[2] Note that specifying numbers in the absence of information runs the risk of falsely suggesting that these exact numbers are known.

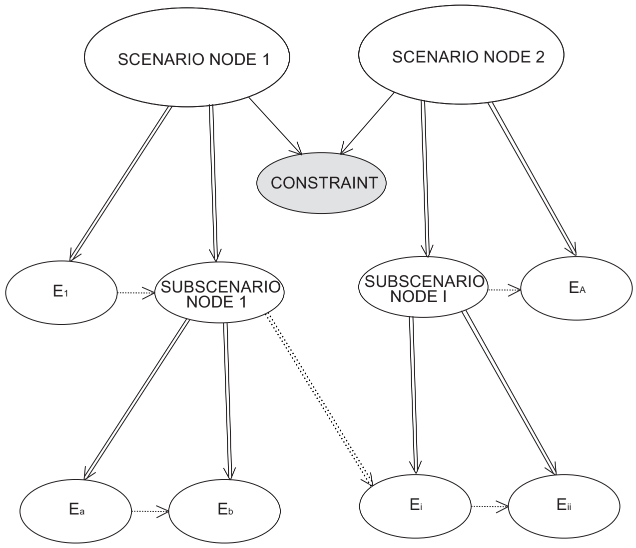

Fig. 8: The merged scenarios idiom

alternative scenarios could be the following: (1) A and B committed the burglary, (2) A visited on an unrelated occasion, explaining the fingerprints and (3) B committed the burglary alone. Note that scenarios (2) and (3) can be the case simultaneously. In this situation, the merged scenarios idiom should only be applied to scenarios (1) and (3). Scenario (2) can still be modeled in the network as an alternative explanation of the evidence, but will not be connected to the constraint node.

With the merged scenarios idiom, two or more representations of scenarios are merged. A constraint node is connected to all scenario nodes as shown in Figure 8: for any scenario node S_i in the collection of n scenarios $\{S_1, \ldots S_n\}$ there is an arrow from S_i to the constraint node C. The constraint node has values s_i for each scenario s_i and one value NA (for not applicable) to denote that an illegal combination of nodes is considered. The CPT will be such that unless exactly one scenario holds, the constraint node will have value NA. Evidence for the constraint node is now set such that the prior probabilities of the scenario nodes behave as desired (see Fenton et al (2011)), while setting the probability of value NA to 0 will make sure that multiple scenarios cannot be true simultaneously, nor can it be that none of the scenarios are true.

Note that this constraint node differs from the standard solution for enforcing mutual exclusiveness as described in Jensen and Nielsen (2007). It has been shown by Fenton, Neil and Lagnado (2011) that the solution in Jensen and Nielsen (2007) captures prior probabilities accurately only when they are distributed uniformly. We therefore suggest to use the solution as proposed by Fenton et al.

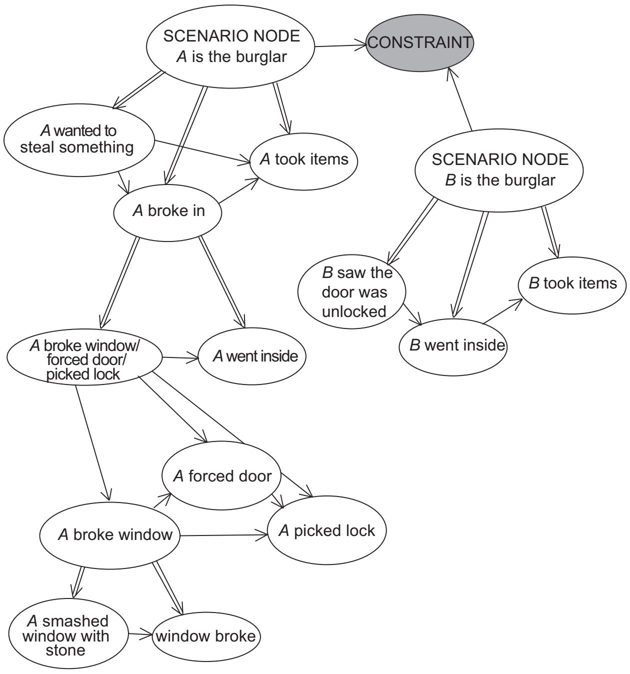

Fig. 9: The scenarios for the burglary case merged

The two example scenarios from the burglary scenario are merged with the merged scenarios idiom in Figure 9.

After using the merged scenarios idiom to put a constraint on the scenarios, there can be overlapping elements: there may be elements that occur in multiple scenarios. For example: in a third burglary scenario, there could be another sequence of events leading to the window breaking. The event 'window broke' will then occur in both scenarios. When this happens, it will be represented only once and any parents of the original nodes are now parents of the new node: arrows from both scenario nodes will point to this one representation of that event or subscenario. The probability table of the new node will require some additional probabilities to be elicited, such as how the parents interact as they influence the new node. This can be done using existing elicitation methods.

As a final step after merging all scenarios for a case, evidential nodes are included in the graph using evidence idioms from (Fenton et al, 2013). These idioms provide the local structures for modeling, for example, the accuracy of a piece of evidence, or the credibility of an alibi as related to a suspect's guilt. Examples of the use of these idioms can be found in our case study in Section 4.

3.2 Unfolding a scenario

A legal case may require very specific details about one part of the case, and much less detail about other parts. In the burglary case the exact events surrounding the break-in itself deserve a lot of attention, including, for example, the exact order in which the burglar entered different rooms in the house. However, the burglar could probably go into a lot of detail about why he decided to break in. A judge or jury will not want to hear all the details about this, but they do want to hear many details about the actual break-in. Legal cases thus require a scenario about the crime that has different levels of detail for the various events. Narratives have the property that they can be told at various levels of detail. In daily life, we employ this property often; for example, when telling a friend about our restaurant visit last night, we would include much detail about the food, and much less about the events of asking for the check, paying and leaving (Schank and Abelson, 1977). However, if we were to report to the police about this restaurant visit because some money was stolen, we would focus on the payment rather than the quality of the food. Wagenaar, Van Koppen and Crombag argue in their Anchored Narratives Theory (Wagenaar et al, 1993) that a story about a crime should be made more precise whenever more details are needed.

With the method of unfolding a scenario, first discussed in (Vlek et al, 2013c), a Bayesian network for a case is gradually constructed, employing the narrative property that various parts of a scenario can be elaborated upon at various levels of detail. Starting with an initial scenario about the case, questions are asked to determine step by step whether more detail is needed about elements of this scenario. When more details are required, the element node from the initial scenario is *unfolded* to a subscenario. An initial version for one scenario from the burglar case was shown as an example in Figure 3. Since the event '*A* broke in' required more details, this event was unfolded to a subscenario describing it in more detail (see Figure 5).

Unfolding a scenario element to a subscenario is done by replacing the element node with an instance of the subscenario idiom. The element node is transformed to a subscenario node, all arrows that were connected to this element node remain connected to the subscenario node, and the conditional probability table for this node remains the same. For example, the event '*A* broke window' was part of the scenario, so the subscenario describing exactly how *A* broke in remains a part of that larger scenario. The details of the subscenario are modeled as element nodes connected to the subscenario node, as in the subscenario idiom. Essentially, unfolding an element node thus means that the node now serves as a subscenario node and the specifics of the subscenario are added as children of this node in the graph.

As said, the core idea of unfolding scenarios is that some elements of a scenario require more detail, but not all. Whether an element of a scenario requires unfolding depends on its connections to (possible) evidence and whether adding more detail makes these connections more insightful. For example, the event '*A* broke in' could be connected to the evidence that there is a broken window. However, unfolding this event leads to a subscenario with an element 'A smashed the window', to which the evidence of the broken window can be connected instead, representing a more direct connection.

Therefore, when modeling a case, the following three questions can be asked to determine whether an element of a scenario requires unfolding:

1. Is there evidence that can be connected directly to the element node? If so, no unfolding is required.
2. Is there relevant evidence for details of a subscenario for this element? If so, unfolding is required.
3. Would it be possible to find relevant evidence for details of the subscenario for this element? If so, unfolding is required.

By asking the three question above, elements of a scenario are unfolded whenever they lead to relevant evidence being included in the network, or if there is no evidence yet, the possibility of more evidence being found (question 3).

Upon unfolding a scenario, probability tables can be specified during the process. For example, the probabilities of the initial scenario can be specified as soon as this initial scenario is put into a scenario structure with the appropriate idiom(s). This is because unfolding will never add incoming arrows to previously constructed instances of idioms. So no new parents will be added to nodes that are already in the network and the previously constructed probability tables will never have to be changed. The only case in which a probability table might need updating, is when two or more scenarios are merged with the merged scenario idiom (see Section 3.1.4). Of course, one is free to choose to wait with filling in all probability tables until after modeling the entire structure.

3.3 The design method in four steps

In Sections 3.1 and 3.2, concepts that can facilitate the process of building a Bayesian network for a legal case were discussed. The concept of unfolding is used to incrementally construct a more and more detailed Bayesian network. The scenarios and subscenarios that are encountered while unfolding a scenario are represented in the network using the narrative idioms. Currently, our methods assume that a number of relevant scenarios are known, such that the problem that is being addressed by the design method is how a Bayesian network can be constructed on the basis of these scenarios. Therefore, the design method consists of the following four steps:

1. Collect: gather relevant scenarios for the case;
2. Unfold: for each scenario, model an initial scenario with the scenario idiom. Then unfold this scenario by repeatedly asking the three questions:
   (a) Is there evidence that can be connected directly to the element node? If so, no unfolding is required.
   (b) Is there relevant evidence for details of a subscenario for this element? If so, unfolding is required.

(c) Would it be possible to find relevant evidence for details of the subscenario for this element? If so, unfolding is required.

Use the subscenario idiom to model the unfolding subscenarios and the variation idiom whenever a variation is encountered. The process of unfolding is finished when the three questions indicate that no more relevant evidence can be added to the structure;

3. Merge: use the merged scenarios idiom to merge the scenario structures constructed in the previous step;
4. Include evidence: for each piece of evidence that is available, include a node and connect it to the element node it supports. Additionally, include nodes for evidential data that is to be expected as an effect of elements in the structure.

4 Case study

In this section, the design method discussed in Section 3 is used to model the murder of Leo de Jager. This murder case is part of the so-called ‘Anjum case’. Marjan van der E., owner of the boarding house (in the village of Anjum) where Leo was killed, was convicted. Later, the case was re-investigated by legal scholars in a project called ‘Project Gerede Twijfel’ (Project Reasonable Doubt).[3] In this project, conducted by scholars from the VU University Amsterdam and the University of Maastricht, criminal cases with a definitive conviction are investigated “if there is a real possibility that an innocent person was convicted”.

The case study presented here is based on two sources: the book by Hans Crombag and Han Israëls (2008) about their investigation of the Anjum case for the Project Gerede Twijfel and the analysis by Floris Bex (2011), which was in turn based on the book by Crombag and Israëls. Since these two books were our only source of information for this case, our results are undoubtedly influenced by the ideas presented by Crombag, Israëls and Bex. The aim of this case study is therefore not to evaluate the case objectively, but rather to evaluate our techniques by modeling this complex case in a Bayesian network.

Crombag and Israëls (2008) formulate four scenarios concerning the murder of Leo de Jager. In order to model a complete network for a case, the design method of course requires all relevant scenarios to be modeled. In this case study, the modeling of two of the four scenarios is discussed. In Section 4.7 the results of the network are discussed, and the conclusions that can be drawn from it. Section 5 then evaluates our techniques and how well they were capable of modeling this case.

Both (Crombag and Israëls, 2008) and (Bex, 2011, Chapter 6) use the same fictitious names for most persons involved in the case. Only for the prime suspect, Marjan van der E., her real name was used. We will use the fictitious names as they were introduced by Crombag and Israëls.

[3] www.projectgeredetwijfel.nl (in Dutch)

4.1 The case

On the evening of December 24th, 1997, Evert Beekman came into the police station to report a murder. Beekman said he had seen a dead body on the property of Marjan van der E., and that he had recognized the body as Leo de Jager. Furthermore, on Beekman's instructions, the police dug up the remains of another body in Marjan's garden, recognized as Herre Sturmans. It is the murder of Leo de Jager that will be discussed here.

4.1.1 The people involved

Marjan was the proprietor of a boarding house in Anjum. Leo rented a small house from Marjan in Moddergat and did some odd jobs around the boarding house. Beekman was a dealer in timber in Anjum and he knew Marjan because she had placed orders with him in the past. Other important persons involved in the case are Marga Waanders, who was staying in the boarding house at the time of the murder, and Eef Tasman, who did some administrative work for Marjan occasionally. Finally, Jaap Kuilstra had heard from Beekman about the murder, advised him to go to the police and came in with him to the station.

As it turned out later, Marjan, Beekman and Kuilstra had a cannabis operation in Marjan's barn. At the time that Beekman reported the murder, the police had found the operation a week before, and Marjan was a suspect. However, she had denied any involvement and claimed that she had let the barn to someone else. She had then promised that she would show the police a contract of this agreement.

4.1.2 The evening of December 24th

Initially, Beekman reported that Marjan came to him to tell him that she had killed someone. This was around 7 in the evening. She returned to the boarding house and Beekman arrived there soon after. He talked to Waanders for a while, while Marjan was cleaning the hallway: she said that Leo had puked there and they weren't allowed to see. In all, Beekman and Waanders were talking for about 10 minutes. At some point, Waanders went into the hallway to get a wash cloth and returned shortly after. Beekman later went into the hallway and saw Marjan scrubbing the floor. Beekman saw blood in the hallway and recognized a trail of blood which he thought might be from the back of a head as a body was dragged to the front door.

Then Marjan took Beekman to the front door, where he saw a dead body lying outside under tent canvas and recognized it as Leo's. Beekman reports that the victim's head was injured in 6 or 7 separate places, which he assumed was the result of hitting the head with a sharp object.

Later, Beekman said that Marjan did not just tell him that she had killed someone, but that she had killed Leo. Furthermore, he admitted that he helped Marjan to wrap the body in the piece of canvas. He then returned to the boarding house at 2 a.m. to help Marjan drag the body to the front yard. Kuilstra confirms this story and explains that he had advised Beekman not to tell this part to the police.

When Marjan was first interrogated, she seemed too confused to say anything informative. In any case, throughout the investigative process she persistently claimed that she did not kill Leo, nor did she drug him. Waanders, who was at first also a suspect, did give a statement right away. She said that Leo was at the boarding house when she arrived in the afternoon. She last saw him in the hallway around 6 p.m., talking to Marjan who was trying to convince him to stay in one of the rooms in the boarding house. In later interrogations, Waanders gave some more details about the events of December 23rd and 24th. In the afternoon, Leo seemed under the influence of something. Waanders also mentioned that she found him in the barn at some point, and took him back to the house. Marjan then seemed agitated to find Leo back in the house. Marjan gave Leo a glass of warm water with jenever (Dutch gin), which she called 'a grog'. Later, when Marjan did not show up for dinner, Waanders took a look in the hallway and saw Marjan giving Leo another glass of jenever. Later, Waanders also said that she had 'images' of Marjan hitting Leo, but she said that these images do not mean that she actually saw this.

Waanders explicitly stated that she did not see any blood stains in the hallway. She did see Beekman when he came over, around dinner time. She then went to get a wash cloth because her eyes were irritated. Later in the evening, Waanders saw two shadows standing outside, possibly Beekman and Marjan. She also saw someone digging a hole in the front yard some time on the 23rd or 24th of December. At the end of the evening, Marjan and Waanders had a drink together, and went for a walk with their dogs, Waanders said to the police.

4.1.3 The evidence

Statements made by Beekman, Waanders, Marjan and Kuilstra serve as evidence in this case. Their main points were summarized in the description above. Additional information will be discussed as soon as it is of interest for the construction of the Bayesian network below. In this section, we present the key evidence other than the testimonies.

When investigating the boarding house on December 25th, the police found traces of blood in several places. Most of the blood traces were in the hallway. Furthermore, a wad with a bloody knot of hair was found in the trash can in Waanders' room. The police also found two hammers, a large one and a regular sized one, with watery bloodstains on them. These hammers were found in the barn. For each of these blood stains, a DNA match was found with Leo, though the profiles drawn from the analyzed material were not complete. For the blood in the hallway the probability that this was from a random other person than Leo was estimated to be much less than 1 in a million. For the blood on a hammer, this probability was 1 in 100 for the hammer head, and 1 in 1700

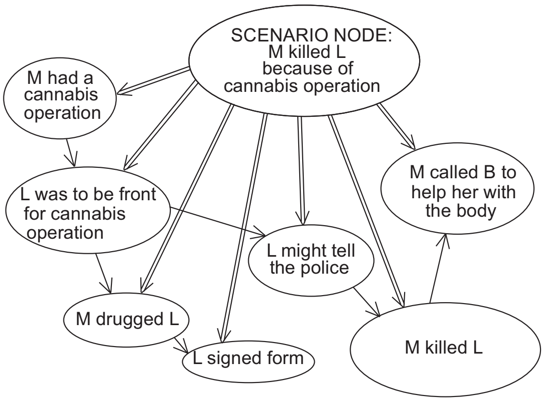

Fig. 10: An initial scenario for the Anjum case

for the hammer handle (although the latter estimations have been disputed by another expert).

In the trash in the kitchen three empty strips of the medicine Temazepam were found, and a strip with ten empty capsules that were cut open. Additional empty capsules were found in the trash, plus a medicine bottle in the name of Leo de Jager. An autopsy of Leo's body showed high concentrations of Temazepam and alcohol in Leo's blood. The level of Temazepam was far more than the amount advised for daily use, and in fact far above the toxicity level. Temazepam is not lethal, but it will cause some strange behavior A pathologist concluded that the cause of death must have been a heavy blow to the head, leading to a fractured skull.

Finally, part of the evidence that was used in the trial concerning the Anjum case had to do with a bank fraud in which Marjan was supposedly involved. This may have been an additional motive for killing Leo, as Marjan seemed to be working towards transferring money away from Leo's account. To keep our model compact, we have chosen to leave out this motive and the evidence related to it.

4.2 Scenario 1: Marjan killed Leo

The first scenario that is modeled here (Crombag and Israëls' second scenario) is an elaborate version of what can be distilled from the police investigation. In this scenario, Marjan killed Leo and Beekman helped her move the body.

4.2.1 An initial scenario

The initial first scenario goes as follows:

> Marjan had a cannabis operation. She wanted to use Leo as a front for this cannabis operation. She drugged him because she wanted him to sign a contract. Leo signed the contract. Marjan was worried that Leo might tell the police, so she killed him. After this, she went to call Beekman, who helped her to drag the body to the front yard.

With the scenario idiom, this is modeled as shown in Figure 10. There are a number of connections between variables within the scenario. For example, because Marjan had a cannabis operation, she planned to use Leo as a front for this operation (hence M `had a cannabis operation` has an arrow to L `was to be front for cannabis operation`). Similarly, this plan to use Leo as a front led to Marjan drugging him, which in turn resulted in Leo signing a form (L `was to be front for cannabis operation` has an arrow to M `drugged` L which has an arrow to L `signed form`). However, as a result of this plan, Marjan also worried that Leo might tell the police, which resulted in Marjan killing Leo (arrows from L `was to be front for cannabis operation` to L `might tell police` and from there to M `killed` L). Finally, because Marjan now had to deal with the results of killing Leo, she asked Beekman for help (arrow from M `killed` L to M `asked` B `to help her with the body`).

For each of the nodes in this structure, there is possibly a subscenario to unfold. To determine whether a node should be unfolded, the three questions from Section 3.2 are asked. In the sections that follow, the unfolding of specific nodes is discussed. Figure 18 in the appendix shows the resulting structure after unfolding.

4.2.2 The cannabis operation

The leftmost node 'M had a cannabis operation' has some evidence that can be connected to it directly (answering question 1 with yes): a police report from another investigation in which the police found a cannabis operation in her barn. Therefore, unfolding is not required. Of course it is still possible to unfold this node: the police report of the cannabis case surely relies on more detailed evidence about the cannabis operation. However, in order to keep the graph from getting too complex, those details are left out of this particular murder case.

4.2.3 Marjan's motive

The next node in the initial scenario, 'L was to be front for cannabis operation', has no evidence that can be connected to it directly (question 1: no). But there are some indications (question 2: yes) that Marjan was indeed planning to use Leo as a front for the operation: her accountant Tasman testified that he made up a false contract, and this false contract was found in Marjan's house. By unfolding this node, the relation of this evidence with the state that L was to be a front for the cannabis operation can be specified.

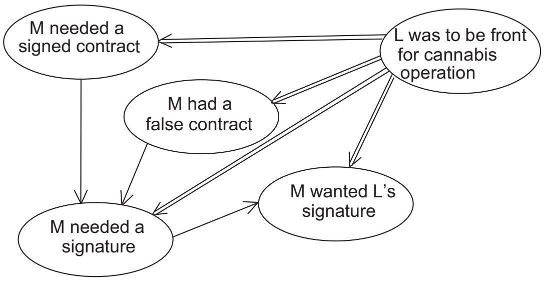

Fig. 11: A subscenario for Marjan's motive

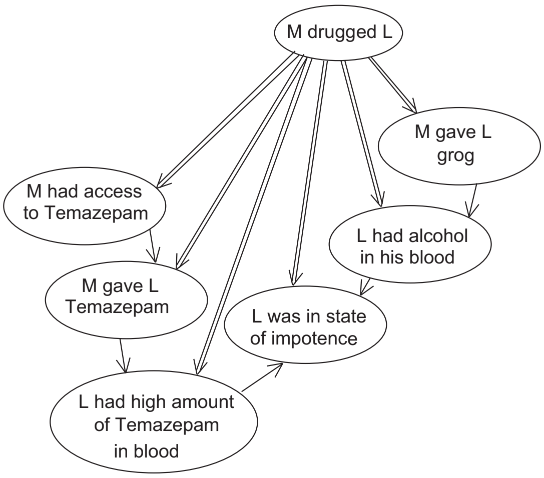

Fig. 12: The drugging of Leo

This results in a subscenario structure as shown in Figure 11. Marjan needed to present a contract to the police to support her alibi that she had rented out her barn to someone else. She had a false contract made and she just needed a signature. These variables together resulted in Marjan needing someone's signature, which in turn resulted in her wanting to get Leo's signature. The aforementioned evidence will be included as a final step in the design method (after representing and merging the scenarios).

4.2.4 Marjan drugged Leo

The node 'M drugged L' certainly requires unfolding: there is no evidence for it directly (question 1: no) but there is evidence that Leo was drugged (Temazepam in his blood) and that Marjan had access to these drugs (bottles of Temazepam, question 2: yes).

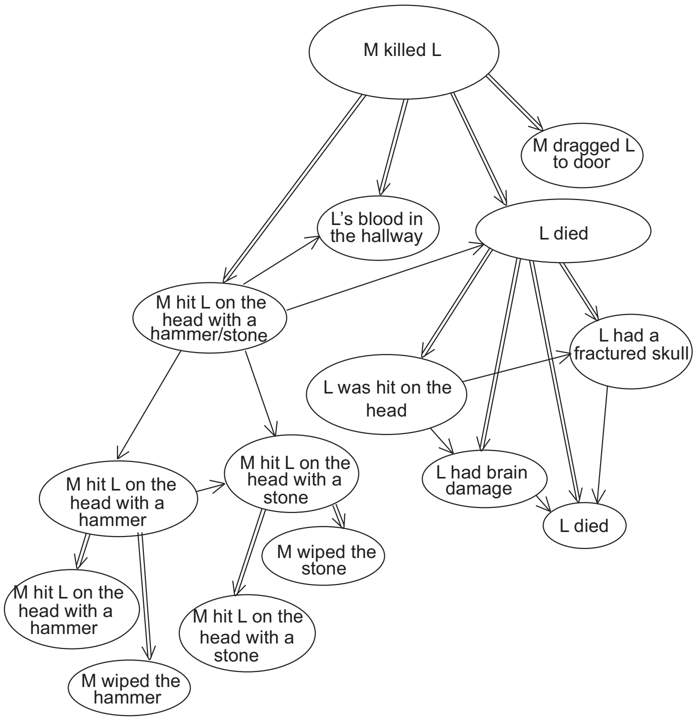

Fig. 13: A subscenario for Marjan killing Leo

This subscenario is shown in Figure 12. Marjan had access to Temazepam, which she then gave to Leo, resulting in Leo having a high amount of Temazepam in his blood. Moreover, she also gave him grog, which led to alcohol being present in Leo's blood. As a result of the high amount of Temazepam and the alcohol in his blood, Leo was in a state of impotence.

4.2.5 Leo signed form and Leo might tell the police

The nodes 'L signed form' and 'L might tell the police' do not need to be unfolded. In the first case, this is because there is direct evidence (question 1: yes) because a form with Leo's signature was found and further unfolding the node will probably not lead to any more relevant evidence (questions 2 and 3: no). As for the node 'L might tell the police', there is no evidence for it directly (question 1: no) but it is also not likely that any relevant evidence will turn up by unfolding the node (questions 2 and 3: no). Therefore, the node is left as it is.

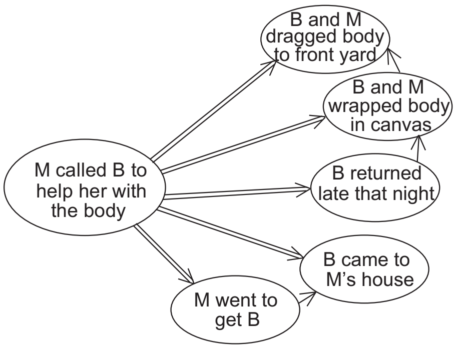

Fig. 14: A subscenario for Marjan asking Beekman for help

4.2.6 Marjan killed Leo

The node 'M killed L' has no direct evidence (question 1: no) and more evidence can be included by unfolding it. In fact, several pieces of evidence can be included: there is blood on the hammer, there is blood in the hallway, a wad with blood and hair on it was found and there is an autopsy report about Leo's death (question 2: yes). According to a testimony from Beekman there was a bloody trail as a result of Marjan dragging Leo to the door. Finally, there is Marjan's testimony, where she denies that she killed Leo.

See Figure 13 for the resulting subscenario structure. In this subscenario, it is specified that Marjan hit Leo in the head with either a hammer or a stone using the variation idiom. Within the variation idiom, two subscenarios are included in which Marjan wiped the stone (in the stone variation) or the hammer (in the hammer variation) to remove Leo's blood. As a result of Marjan hitting Leo with either a hammer or a stone, Leo's blood ended up on the floor of the hallway and Leo died. Subsequently, Marjan dragged Leo's body to the door.

The node that Leo died can be unfolded to include more detail: as a result of being hit on the head, Leo had a fractured skull and brain damage, resulting in his death.

4.2.7 Subscenario: Marjan went to get Beekman

Finally, there is the node 'M called B to help her with the body'. There is evidence that Marjan went to Beekman, but this is not direct evidence for the node (question 1: no) but only for part of the subscenario that can be unfolded here (question 2: yes). Furthermore, there is evidence that Leo's body was dragged to the front yard (Leo's body and a trail of his feet being dragged).

This amounts to a subscenario structure as shown in Figure 14. Marjan came to Beekman to tell him that she killed Leo, and as a result he came with

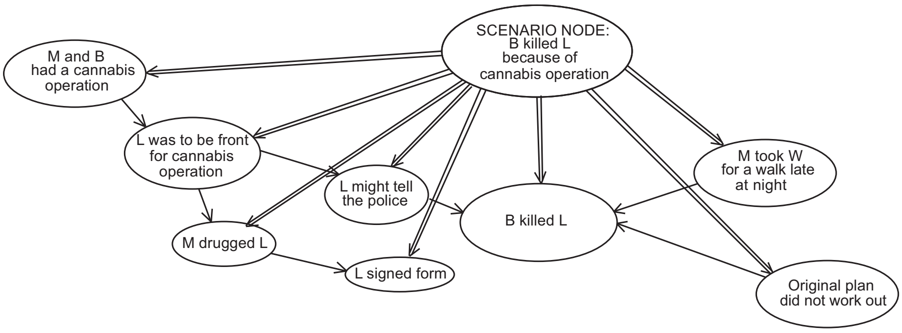

Fig. 15: The second scenario

her to the house. Later that evening he returned, to help her wrap the body in canvas and drag the body to the front yard.

4.3 Scenario 2: Beekman killed Leo

In an alternative scenario, Beekman killed Leo with Marjan's help. This is based on Crombag and Israël's third scenario (Crombag and Israëls, 2008). A number of elements are the same as in the previous scenario: once again the motive is to use Leo as a front for the cannabis operation. Again Marjan drugged Leo because she wanted him to sign the form. However, in this scenario the original plan was that Marjan would take Leo to the barn, such that Beekman could kill him. However, this plan did not work out properly because Marga Waanders took Leo back to the house. This resulted in a new sequence of events in which Beekman killed Leo late at night while Marjan and Waanders went for a walk with their dogs. This is modeled with the scenario idiom in Figure 15.

Connections between variables within the scenario are similar to the connections in the initial scenario for Marjan killing Leo: this time the cannabis operation by Marjan and Beekman together made them want to use Leo as a front. Because of this goal, Marjan drugged Leo which led to him signing the form, but another result was that Leo might tell the police. Together with the original plan not working out and the late night walk by Marjan and Waanders, this led to Beekman killing Leo (so there are arrows from L `might tell police`, from `Original plan did not work out` and from M `took W for a walk late at night` to B `killed` L).

For this scenario, several subscenario structures from the previous scenario can be reused, namely: 'L was to be front for cannabis operation', 'M drugged L' and 'L died because he was hit on the head with an angular object'. In the following sections, the other nodes will be discussed and after the necessary unfolding, this results in a scenario structure as shown in Figure 19 in the appendix.

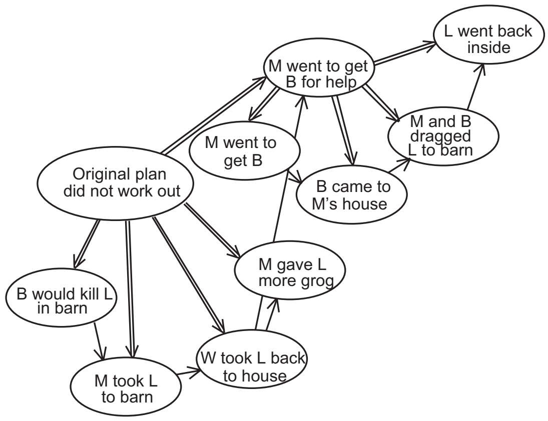

Fig. 16: A subscenario for the original plan not working out

4.3.1 The original plan did not work out

The node 'Original plan did not work out' has no evidence that can be connected to it directly (question 1: no), but it requires unfolding such that Waanders' testimony can be included as evidence (question 2: yes): she somehow found Leo in the barn and brought him back to the house, which seemed to upset Marjan. The subscenario is modeled in Figure 16.

Originally, Beekman would kill Leo in the barn, so Marjan took Leo there. But because Waanders then took Leo back to the house, Marjan gave Leo more grog and went to get Beekman for help. The latter can be unfolded to a subscenario: Marjan went to get Beekman, so Beekman came to her house and together they dragged Leo to the barn. However, Leo now went back inside by himself.

4.3.2 Marjan took Waanders for a walk

The node 'M took W for a walk around midnight' does not need to be unfolded: there is direct evidence (question 1: yes), in the form of Waanders' testimony.

4.3.3 Beekman killed Leo

Finally, the node 'B killed L' has no evidence to be connected directly (question 1: no), so it needs to be unfolded such that evidence about the blood on the hammer, in the hallway and on a wad can be included, as well as statements from the autopsy report (question 2: yes). The subscenario is shown in Figure 17.

Leo was now in the house, and Beekman hit him on the head with either a hammer or a stone, resulting in Leo's death and Leo's blood in the hallway. The variation idiom is again used to model the hammer/stone variation, and the

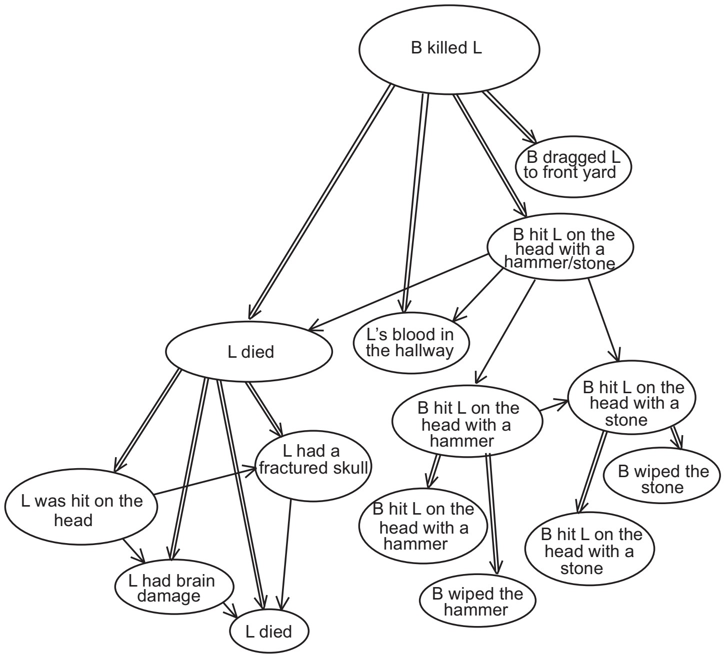

Fig. 17: A subscenario Beekman killing Leo with a hammer

subscenario for Leo's death is the same as before. Finally, Beekman dragged Leo to the front yard.

4.4 Merging the two scenarios

For merging the two scenario structures in Figure 18 and 19, the merged scenarios idiom is used. The resulting structure is shown in Figure 20 in the appendix.

A constraint is put on the two scenario nodes to make sure that at most one of the two scenarios is true, but not both: either Marjan killed Leo, or Beekman killed Leo. There is quite some overlap in the two scenarios. Some subscenarios occur in both scenario idioms. Furthermore, within the subscenarios there are some overlapping nodes, such as "the hammer was wiped" and "*L*'s blood in the hallway". For such cases, one such subscenario or node is kept. For clarity, the nodes that are part of the second scenario only are dark gray Those that are part of the first scenario only are white, and those that are in both scenarios are light gray

4.5 Including the evidence

After merging the scenarios, the evidential nodes are included. This comprises evidence about police findings, witness testimonies and forensic reports such as a toxicology report and an autopsy report. See Figure 21 in the appendix for the resulting structure.

Each piece of evidence is modeled with the evidence idioms from (Fenton et al, 2013), that is: an evidential node accompanied by a node 'accuracy of evidence', abbreviated to `acc`. The evidential node is connected to a node in the merged scenarios idiom for which it is evidence. For example: the empty strips of Temazepam (evidence) are evidence for the hypothesis that Marjan gave Leo Temazepam (a node in the scenario). The accuracy of evidence node captures the possibility that a piece of evidence is not correct: there may have been a mistake in the lab (when it comes to DNA tests, for example) and a witness may lie.

Note that the accuracy of each witness was captured as a single node connected to all statements made by that witness: there are nodes `Accuracy of Beekman`, `Accuracy of Waanders`, `Accuracy of Tasman` and `Accuracy of Marjan`. For both Beekman and Marjan the alibi evidence idiom from Fenton and Neil (2000) was used to capture that their guilt in this case would influence their accuracy. According to this idiom, an arrow is drawn from the hypothesis about a witness' guilt (in our case, the corresponding scenario node) to the accuracy of this witness. As a result, if the scenario node for Marjan killing Leo holds, the accuracy of Marjan's testimony is low and if the scenario node for Beekman holds, the accuracy of Beekman's testimony is low.

4.6 Probabilities

This section is concerned with the probabilities for the network constructed in the previous sections (see Figure 21 for the graphical structure). Some, but not all, probabilities in the network will be discussed below to give an impression of the quantitative part of the Bayesian network. Our full quantification of the network can be found in our GeNIe model of this case, available for download at www.charlottevlek.nl.

Several numbers in the probability tables of the network are fixed because of the idioms that were used. For example, in the scenario and subscenario idioms, a number of connections are drawn as double arrows, signifying that a node is an element of the scenario or subscenario. For these nodes, part of the probability table expresses logical relations (see Table 1): an element of a scenario is always true when the scenario as a whole is true. This means that, for example, for the node `L signed form`, the probability table is as in Table 6. Some probabilities remain to be determined: these are the probabilities for Leo signing the form outside the context of this scenario, one for the case where he was drugged by Marjan, and one for the case that he was not. For the purposes of this example the probability that he did sign without the

Table 6: Probability table for L `signed form`

	`Scenario Node = true`		`Scenario Node = false`	
	M `drugged` L `= true`	M `drugged` L `= false`	M `drugged` L `= true`	M `drugged` L `= false`
L `signed form = T`	1	1	0.6	0.1
L `signed form = F`	0	0	0.4	0.9

scenario node being true is estimated to 0.9 when he was drugged, and 0.3 when he was not drugged.

Other probabilities are in some cases available within the evidence. For example, in the subscenario M `drugged` L, there is a node L `was in a state of impotence`. In the toxicology report, an expert stated that given that Leo had alcohol and high amounts of Temazepam in his blood, it was to be expected that he would be in a state of impotence. This means that the probability $P(L$ `was in state of impotence = true`$|$ M `drugged` L `= false`, L `had alcohol in his blood = true`, L `had high amount of Temazepam in his blood = true`$)$ is high, say 0.99. Similarly, the probabilities within the subscenario about Leo's death follow largely from the autopsy report, such as for the node L `died`, connected to L `had a fractured skull` and L `had brain damage`.

As for the accuracy of evidence, forensic evidence is often accompanied with numbers reporting the accuracy of the lab. The accuracy of other evidential data may be more difficult to estimate. For example, Waanders' testimonies have changed a lot during the investigation, so they have been estimated to be less reliable.

Other probability tables for the element nodes can be filled in with the help of elicitation methods, such as for example a verbal-numerical probability scale (Renooij, 2001): on a horizontal or vertical line with numerical and qualitative anchors on either side of the line one indicates the value of the sought probability. In the Anjum case, the event that Marjan killed Leo given that she is not worried about him telling the police and the scenario node is false, can be qualified as 'very unlikely', translated to a probability of 0.001 in the network. The probability that Marjan had a cannabis operation (without the scenario node being true) is unlikely (0.05), while the probability that Leo signed the form while he was drugged is quite likely (0.9).

In order to evaluate how probable various hypotheses are, posterior probabilities are in the end what matters. In addition to likelihoods (partially) supplied by experts, prior probabilities are needed for the computation of posterior probabilities. These priors are, as in any Bayesian approach, difficult to elicit. If needed, a sensitivity analysis can be used to reveal how various estimations of priors may affect the overall outcome, and also how various conditional probabilities may affect the outcome (see below for an example of a sensitivity analysis). In our network, the priors that need to be specified are the numbers in the probability tables for the scenario nodes. They represent the prior probability of a scenario, which could arguably be viewed as the plausibility of a scenario as it occurs in narrative research, expressing how likely a person would find a scenario beforehand. One can really only subjectively estimate these probabilities. Any attempt to objectively estimate this plausibility gives rise to a number of issues. Firstly, there is the principle that in court no prejudice should be held against any suspect, which is sometimes argued to mean that prior probabilities should be equal for everyone. However, this does not imply that two scenarios about different suspects should always have equal prior probabilities: that also depends on the coherence of the scenario. Secondly, there is the reference class problem (Handfield, 2012), concerning how much detail to include in the prior probability. Knowing that a suspect lived in the neighborhood where the crime took place undoubtedly leads to a higher prior probability (because there was more of an opportunity) while knowing that a suspect lived on the other side of the world decreases the prior probability. But even when the prior probability does not take into account in what part of the city the suspect lived, there will always be some prior information included, such as that the suspect lived somewhere in the country, on the continent on somewhere in the world.

A subjective estimate of the priors still runs into the principle of equality in court: no distinction should be made based on prejudice. Therefore, in our model of the Anjum case, both prior probabilities for the scenario nodes being true were set to 0.001. Another problem with a subjective estimate, however, is that once the model is done, one might be tempted to view these numbers as if they are objective. In other words: explicitly quantifying the priors might lead to the false illusion of objectivity. We strongly emphasize that our method is meant to formalize subjective accounts of scenarios. In this perspective, our method supplies techniques to support a subjective decision with more formal tools.

Finally, after filling in all probabilities underlying the network, a sensitivity analysis (Jensen and Nielsen, 2007) can be performed on nodes with disputable probability tables. Suppose, for example, that one is unsure whether the probability of Leo signing the form without being drugged (and the scenario node being false) should really be 0.3, as was estimated in the model. One could then think of what values this probability could be instead, say 0.1 the lowest and 0.5 the highest. A sensitivity analysis can then indicate how strongly a change in this value will influence the probability of the scenario node in which Marjan killed Leo. Performing this analysis on the node `Leo signed form`, it shows that the probability for the scenario node does not change (when rounded off to two decimals), whether the value of the said probability is 'high' (0.5), 'middle' (0.3) or 'low' (0.1). This leads to the conclusion that, given the other probabilities as they were filled in for this case, the probability for Leo signing the form without being drugged can be left as it was since it is apparently not very influential. Note that such a sensitivity analysis determines the influence of a change in probabilities given that all other probabilities remain fixed. When multiple nodes in a network have disputable probabilities, a sensitivity analysis on multiple nodes should take into account any configuration of changed probabilities.

4.7 Conclusions drawn from the network

In this section, the posterior probabilities of the scenarios are compared. In order to draw conclusions from the network, our model has been specified in the Bayesian network tool GeNIe 2.0.[4] Our model is available for download at www.charlottevlek.nl. Running our network in GeNIe 2.0 shows that, given all evidence, the scenario that Beekman killed Leo is more likely than Marjan killing Leo.

By entering the evidence piece by piece, the network gives some insight into the case, explicating how this conclusion was reached. The network reveals that the scenario for Marjan being the killer is very likely, based on large subsets of the evidence. But as soon as evidence is entered to disprove some of Beekman's testimonies, the probabilities change to show that Beekman is more likely to be the killer, with some room for doubt. This reflects the conclusions by Crombag and Israëls (2008).

When no evidence is entered in the network, both scenario nodes (Marjan killed Leo and Beekman killed Leo) are very unlikely, as their prior probability was estimated close to 0. All numbers in this section will be rounded off to two decimals, which means that the prior probability of 0.001 for the scenario nodes is now rounded off to 0. Loosely following the police investigation, we start by entering Beekman's testimonies. This leads to a somewhat higher probability for the scenario with Marjan as a killer (0.04), while the scenario about Beekman is still at almost 0 probability. Next Marjan's testimony is inserted in the network, but the numbers do not change: since her accuracy is connected to the scenario node, her testimony is taken to be most probably (0.95) inaccurate.

By adding more evidence to support (parts) of the scenario with Marjan as the killer, this scenario node obtains a higher probability. With the evidence of Marjan being a suspect for the cannabis operation, the probability of the scenario node becomes 0.17 (scenario about Beekman: still 0). Adding evidence about Leo's signature on the form, Tasman's testimony about this form and a false contract being found, the probability of Marjan being the killer becomes 0.73 (versus Beekman: now 0.01). Various pieces of evidence were found to support that Marjan drugged Leo: a prescription and empty strips of Temazepam, a toxicology report and Waander's testimony that Marjan gave Leo a grog. This results in a high (almost 1) probability for the subscenario node `Marjan drugged Leo`, while the scenario node for Marjan killing Leo now becomes 0.93 (Beekman: still 0.01).

Next, a number of forensic pieces of evidence about the killing itself are entered: Leo's blood that was found on the hammer, in the hallway and on a

[4] GeNIe 2.0 is available for free on genie.sis.pitt.edu

wad with blood and hair, his dead body in the yard and a trail of his feet being dragged through the yard, and the autopsy report. These result in a probability of 0.99 for the scenario that Marjan killed Leo, leaving the Beekman scenario at 0.01. Note that at this point no evidence has been included about *who* killed Leo, except for a testimony from Beekman, stating that Marjan said she killed Leo. The evidence that was entered so far all supports the scenario that Marjan did it, so there is no reason to be suspicious about Beekman's testimonies. For a moment we can take out Beekman's testimonies, just to see what happens to the probabilities: this leads to a probability of 0.63 for the scenario about Marjan, and 0.37 for the scenario about Beekman.

Now let us return to the situation in which all of Beekman's testimonies are entered in the network, plus forensic evidence about Leo's death and the evidence about the cannabis operation and the contract. A number of relevant statements were made by Waanders: she had 'images' of Marjan hitting Leo, but she did not see any blood in the hallway. She saw Leo's feet being dragged away through the hallway. She saw Beekman when he came over around dinner and she went for a walk with Marjan and their dogs. Entering all of these testimonies leads to a slightly lower probability for the scenario in which Marjan did it: 0.94, while the scenario about Beekman is now on a probability of 0.05.

Finally, a crucial piece of evidence turns out to be the fact that there was no bloody trail found in the hallway, even though Beekman said in his testimony that he saw such a trail. Taking into account that a scenario in which Marjan killed Leo would lead to such a bloody trail (she supposedly dragged Leo outside the door by herself) and that other blood traces were very obvious and easy to find for the forensic team, this piece of evidence surely contradicts the scenario about Marjan, and it furthermore compromises Beekman's accuracy. Therefore, the updated probabilities are now 0.25 for the scenario about Marjan and 0.74 for the scenario about Beekman.

5 Discussion

Our design method was developed with the following three goals in mind:

1. By taking a narrative approach, the method aims to prevent tunnel vision;
2. By combining the narrative with a more formal, probabilistic approach, the method should avoid the pitfall of a good story pushing out a true story, commonly seen as a negative effect of using narratives;
3. By taking narratives as a basis for the construction of a probabilistic model, the method should provide heuristics for finding the relevant variables during the process of building a Bayesian network.

In the following subsections, each of these three goals are evaluated in the context of our case study.

5.1 Tunnel vision

In the Anjum murder case, one scenario, that of Marjan killing Leo, dominated the case. In a narrative approach, there is a focus on considering alternative scenarios. Though alternative scenarios are not strictly required for reasoning in court, considering various alternatives provides an approach that may help prevent tunnel vision. In our case study in Section 4, the alternative was the scenario about Beekman killing Leo. By always asking for multiple scenarios, our method takes advantage of this property of the narrative approach, thereby reducing the risk of tunnel vision (1).

There is room for improvement in the prevention of tunnel vision, by making our approach more dynamic in that it could better handle changing and expanding scenarios during the investigative process. Currently, our design method assumes that in an earlier phase, a number of scenarios were formulated. The method then results in a Bayesian network modeling these scenarios. However, during an investigation, scenarios about what happened are gradually formed. The same happens for the Anjum case in (Crombag and Israëls, 2008), for example concerning the killing of Leo. For a long time it remains unclear at what time Leo was killed and by whom. Based on various testimonies by Beekman and Waanders, multiple possibilities are taken into account, but only two possibilities deserve further investigation: either Leo was killed by Marjan some time before Beekman came to her house, or Leo was killed by Beekman some time later that evening. These two possibilities are the core of the two scenarios modeled in Section 4, while the more unlikely possibilities never made it into the network. Modeling this process of gradually constructing a theory during the investigation is a topic for future work. Moreover, an extended design method that can help to find multiple scenarios during the process of investigation could actively help to prevent tunnel vision.

5.2 A good story versus the true story

In the conclusions of our network in Section 4.7, the scenario about Beekman killing Leo turned out to be more likely than the scenario about Marjan killing Leo. Of course, there is no way to know which scenario was in fact true, so we cannot be sure that our methods helped to choose the true story over a good story (2). However, in the actual trial for this case, Marjan was convicted for the murder. Only in the thorough investigation by Crombag and Israëls (2008) did the scenario about Beekman turn up as a better explanation of the evidence. Apparently, the scenario in which Marjan killed Leo was good enough to be believed by many people for a long time: it seems that this was the good story that pushed out the true story in the trial. A closer look at the network in Section 4.7 revealed how Marjan being the killer came to be so believable: most of the evidence supported (or at least, did not disprove) this scenario. However, as soon as some evidence was instantiated in the network to disprove one of Beekman's testimonies, the probabilities flipped, showing that Beekman probably lied and most likely committed the murder himself. Whereas a judge or jury might find it difficult to let go of their earlier belief that Marjan was the killer, the network reveals that Beekman's testimonies are likely to be false given all evidence. Formalizing the narrative approach thus helps to choose the more likely story over a good story.

There is one limitation to this formalization: it requires a high level of precision, both for the required numbers and for the structure of the graph. There are elicitation techniques to help find all the numbers and the structure of the graph, but it is not clear how successful these techniques are for this particular application to narratives in the legal field. As was argued in Section 4.6, a Bayesian network is very useful as a representation of a subjective account of the case. It is then a useful tool to structure a subjective thought process and to compare relative probabilities of multiple scenarios. The risk of such a formalization is the illusion of objectivity it creates: an elaborate Bayesian network for a case might appear impressive to someone who was not involved in the construction. Quantifying all probabilities with numbers may lead to the impression that all these numbers are really known: unfortunately, in a Bayesian network there is no way to indicate that a probability is unknown. Our method is explicitly not intended to calculate any objective posterior probability, but rather to help formalize the decision process of a judge or jury. The formalization helps to find errors in reasoning such as a good story pushing out a true story as described above.

5.3 Heuristics for the building process

During the construction of the network for the Anjum case, the heuristics provided by a narrative approach helped to keep track of the building process and find which variables are relevant and should be included in the network (3). By starting with an initial scenario that gradually unfolds, the process is structured. And a crime scenario forms the context for the evidence in a legal case, thereby indicating which variables are relevant and which are not. As Bex (2011) writes, an advantage of the narrative approach is that stories are helpful when reasoning about events for which there is no evidence: such events can be inferred from other circumstances that are supported by evidence via the transfer of evidential support. This is of particular importance for the Anjum case, since there is no direct evidence about who killed Leo, except for Waanders' untrustworthy 'images' of Marjan killing Leo. Circumstantial evidence thus needs to be taken into account to be able to infer anything about the actual killing. However, it is then not always clear which evidence is relevant to the case and which is not. For example, only in the light of the scenario about Beekman returning late at night to kill Leo, it becomes clear that Waanders' testimony about Marjan taking her for a walk with their dogs is relevant. With our method of unfolding, narratives are used as a basis to find the relevant variables for our Bayesian network. Moreover, by using this method and in particular the four narrative idioms, the resulting graph is clearly structured into modules representing scenarios and subscenarios.

To conclude this discussion: the design method addresses all three goals: tunnel vision, good stories pushing out true stories and heuristics for the construction of a Bayesian network. There is room for improvement in the form of a more dynamic design method. A method that helps to gradually construct scenarios during the investigative process could help to actively prevent tunnel vision.

6 Related work

Recently, a number of researchers have been working on legal reasoning with narrative or Bayesian approaches. Fenton et al. (2013) developed legal idioms, Sileno et al. (2012) are working on animating crime scenarios with agents, Keppens (2011) and Timmer et al. (2013) studied the extraction of arguments from a Bayesian network and Bex (2011) developed a hybrid theory for narratives and argumentation in law. Finally, Keppens and Schafer (2006) worked on the modeling of crime scenarios, but not in the narrative sense: in their work, scenarios are 'descriptions of a combination of situations and events' (without the coherence requirement of the narrative approach). To our knowledge, no research has been dedicated to representing legal narratives in a Bayesian network or to building legal Bayesian networks with narratives as a starting point.

The following sections are dedicated to work related to various aspects of our design method: work related to the use of idioms (Section 6.1), the idea of unfolding a scenario (Section 6.2) and the case study, which will be compared to Bex' case study (Section 6.3).

6.1 Narrative idioms

In classical narrative research in the legal field such as (Pennington and Hastie, 1992) and (Wagenaar et al, 1993), a key feature of narratives is said to be the *coherence* of a scenario. This is manifested in the transfer of evidential support: when one element of the scenario becomes more likely, this influences the scenario as a whole. According to Bex (2009), a scenario is coherent when it fits a story scheme. In our narrative idioms, a coherent scenario is captured with the scenario node and the subscenario node. Since the values of these nodes can never be observed, they are always uninstantiated. From the structure of the network it then follows that when two element nodes are connected to the same scenario node or subscenario node, the element nodes are *d-connected* (see Section 2). This enables a flow of information between all elements of a scenario: when one element becomes more likely, the entire scenario (and all of its elements with it) becomes more likely.

The idea to build up problem-specific networks using smaller fragments was introduced by Laskey and Mahoney (1997). Employing ideas from Object

Oriented Bayesian networks, they work with substructures that are combined to form larger networks. For the specific application to legal cases, Hepler, Dawid and Leucari (2004) proposed to develop specific recurrent fragments, that can be reused for different legal networks. They identify typical substructures such as 'identification', 'contradiction', 'conflict', etcetera. Fenton, Neil and Lagnado (2013) (also in Lagnado et al (2013)) went on to develop a list of legal idioms comprising idioms to model specific types of evidence and other local substructures. We extended the work by these authors with four new idioms, specifically designed to represent legal narratives.

The idioms from Fenton, Neil and Lagnado are aimed at modeling the local structure of a Bayesian network. For example, their accuracy of evidence and alibi evidence idioms were used in our case study to find how exactly to incorporate evidential data in the network. Our narrative idioms are intended to help build the global structure of a network. By considering various scenarios, it becomes clear which variables are relevant to a case and should thus be included in the network. Taking multiple scenarios into account furthermore prevents that a Bayesian network will only be a representation of what the modeler believes is the truth: considering more than one scenario helps to prevent tunnel vision.

6.2 Unfolding a scenario

Legal idioms or modules help to structure a Bayesian network. According to Hepler, Dawid and Leucari (2004), "the object oriented approach aids both the creation of the model (by reusing purpose-built or off-the shelf modules) and its presentation and application (by the ability to view or hide details as desired)." With our method of unfolding a scenario, we aimed to further enhance a structured building process by formulating a guideline for gradually building a Bayesian network. Hepler, Dawid and Leucari (2004) touch upon such ideas briefly. In a similar spirit to their idea of recurring substructures, they write that a typical first degree murder case always decomposes into standard modules 'a murder occurred', 'a felony was committed' and 'the accused is the culprit'. However, for including a more detailed account of what happed, they present no heuristics on how to construct the corresponding model. Furthermore, Hepler et al. do not discuss when a network is sufficiently detailed.

More research on the gradual construction of a Bayesian network was done by Van Gosliga and Van de Voorde (2008). They developed the Hypothesis Management Framework (HMF) as a design pattern for a Bayesian network. These HMFs again take a modular approach. In particular, these modules include probabilities explicitly as connecting nodes in the graph. This means that multiple experts can work on one network simultaneously. The incremental aspect of the construction process thus refers to various steps in the construction being delegated to various experts. For example, one expert could work on (part) of the network structure, while another fills in the corresponding numbers subsequently. HMFs are particularly well suited for frequent changes in the network, making them particularly interesting for the legal field, where insights might change during the investigative process. However, a design method or some other specification of what steps to take to complete the network are not discussed in (van Gosliga and van de Voorde, 2008). In other words, both the work by Hepler et al. and that of Van Gosliga and Van de Voorde lack heuristics on how to employ their modules or substructures to construct a Bayesian network for a case.

Our method of unfolding provides such heuristics using narratives, employing idioms to gradually unfold the scenario into more details. A crime scenario forms the context that can help to find which variables are relevant to a case, while the gradual unfolding of a scenario helps to keep track of the building process. Three critical questions serve as a stop condition, or as a guideline as to when enough detail has been added to the network.

6.3 The case study: comparing the results with Bex' case study

In (Bex, 2011), Bex modeled the Anjum case with his hybrid theory of argumentation and narratives. This creates the opportunity to closely compare the two methods: Bex' hybrid theory and our idioms with the method of unfolding. In Section 6.3.2, the advantages and disadvantages of both approaches are discussed. First a short introduction into Bex' methods is given.

6.3.1 Bex' hybrid theory

Bex' hybrid theory (2011; 2009; 2010) combines formal argumentation (inspired by the ASPIC+ framework of Prakken (2010)) with causal reasoning from stories. A story consists of events and generalizations: rules that express how one event leads to another. Arguments are used to reason about the events and generalizations in a story. This makes the hybrid theory particularly well suited for reasoning with implicit generalizations since they will be explicated in the theory.

Bex' argumentation theory contains evidential data and a stock of knowledge which is formalized as generalizations. Furthermore, specific context-dependent (defeasible) reasons can be added for reasoning with evidence, such as 'when a DNA-match with a suspect was found, then the suspect was the donor of the sample from which the match was established'. Arguments are used to support or contradict events and generalizations in the stories. For example, a generalization in a story could be 'when someone needs money, he will break in', connecting two events 'A needed money' and 'A broke in'. One might argue that it is not true that when someone needs money, he will break in, or that this particular suspect would never decide to break in somewhere when he needed money. Alternatively, an argument using evidence about a bank statement can be used to contradict the event that person A needed money. Finally, one can argue that a story does not complete a proper story scheme. A graphical representation shows evidential data and elements of the stories as boxes with arrows between them to represent generalizations within a story or arguments supporting or attacking parts of the story.

6.3.2 The two representations

Comparing just the graphical structures of the two representations, Bex' hybrid theory and our idioms show close resemblance when it comes to the local representation of the stories or scenarios. For example, the representation of the events closely related to Leo's death consists of nearly the exact same structure in both models.

On a higher level, the two representations differ in how they capture the global coherence of (sub)scenarios. Bex aims to capture this coherence using story schemes: in his hybrid theory, one can argue that a sequence of events does or does not complete a story scheme. In our representation the global coherence is modeled with a (sub)scenario node connected to all elements of a (sub)scenario. Unlike Bex' theory, our methods do not incorporate reasoning about the coherence of scenarios within the network. An advantage of our approach is that the coherence of a scenario or subscenario is clearly visible in the network. Additionally, this makes the resulting network from our idioms modular, with groups of nodes clearly recognizable as (sub)scenarios in the case, and (sub)scenario nodes that can be read as brief summaries of the (sub)scenario they represent.

One clear difference in the graphical structures of the two representations is in the direction of arrows. In a Bayesian network, arrows are commonly directed to represent causality: they model that a state or event (e.g. breaking in) is the effect of another state or event (e.g. lack of money), or that a piece of evidence is the consequence of an event. Note that this is a convention rather than a technicality, since a Bayesian network merely represent independence, not causality (Dawid, 2009). In Bex' representation, the arrows within a story, representing generalizations, are directed causally like in our network. But arrows concerning an argument, such as a piece of evidence supporting an event in the story, go the other way. These evidential arrows in Bex' representation are common for the argumentative approach: with the evidence as a premise, an argument results in a hypothesis as a conclusion. Such 'explanation evoking' rules contrast with the 'expectation evoking' rules that are common in a Bayesian analysis (Pearl, 1988).

Moving beyond the graphical structure of the two representations, there is one evident difference: with the use of Bayesian networks a quantitative approach is taken, while Bex has a qualitative perspective. The disadvantage of using a quantitative approach as done with Bayesian networks, is that all these numbers need to be specified. However, once they are, this probabilistic approach has the advantage that the appropriate weighing of evidential support follows naturally. By describing the case in terms of (conditional) probabilities, one can calculate how the probability of various scenarios change as a result of instantiating evidence. For example, in the Anjum case we specified in the probability tables that Tasman was a reliable witness, while Waanders was less reliable since she often changed her statements. As a result, the effect of their testimonies on the overall probability of the scenario that Marjan killed Leo is not the same (although this also depends on other factors). For his qualitative approach, Bex describes a number of different methods to measure evidential support. One of these methods is to simply count the pieces of evidence supporting each scenario and compare which scenario has the most pieces of evidence. In the aforementioned example, this would lead to equal evidential support for a testimony by Tasman and a testimony by Waanders.

Another difference between the two representations is the use of generalizations within a story or scenario. Bex' hybrid theory makes these generalizations explicit. For example, there is a generalization 'when someone has high amounts of Temazepam and alcohol in their blood, they will be in a state of impotence' between the nodes 'L had high amount of Temazepam in blood' and 'L had alcohol in his blood' to 'L was in a state of impotence'. By making these explicit, arguments for or against such generalizations can be formulated. In our approach, the arrows between nodes in the network can be viewed as the instantiation of generalizations: the arrows between the aforementioned nodes now represent 'when Leo has high amounts of Temazepam and alcohol in his blood, he will be in a state of impotence'. Probability tables then express how strong these connections are, but our methods do not incorporate reasoning about such instantiations of generalizations (arrows in the network cannot be supported with evidence, unlike in the hybrid theory).

There is a fundamental difference between Bex' hybrid theory and our approach in the methods themselves. The hybrid theory switches between narratives and argumentation depending on the circumstances: when working with evidence, an argumentative approach is taken, while explanations of events ask for a narrative approach. In our methods, the narrative and probabilistic approaches are integrated: a scenario is modeled within the framework of a Bayesian network, while the Bayesian network is constructed taking the scenario as a basis.

Finally, in contrast with Bex' qualitative account of the case, our quantitative representation allows us to draw a conclusion about which scenario is most likely. As discussed in Section 4.7, the network showed that the scenario in which Beekman killed Leo is more likely that the scenario in which Marjan killed Leo, but only after adding the evidence that no traces were found of Leo being dragged through the hallway by Marjan. Of course, many numbers used in the quantification of this case are subjective estimates, following common practice in Bayesian modeling. As a consequence, the conclusion about the most likely scenario depends on these estimates.

7 Conclusion

In this paper, a design method for constructing a Bayesian network was presented. The design method was then evaluated with a case study. The design method employs four narrative idioms, which enable the representation of narratives in a Bayesian network. Furthermore, the method of unfolding a scenario is used to gradually construct a Bayesian network for a legal case.

By using an integrated approach of narrative and probability, our design method aims at enhancing the process of reasoning with evidence in various aspects: (1) by always taking multiple scenarios into account, tunnel vision should be prevented (2) formalizing the narrative approach by representing scenarios in a Bayesian network should resolve the problem of a good story pushing out a true story and (3) taking narratives as a basis for the probabilistic approach provides the heuristics for finding the relevant variables for a model.

As the case study showed, multiple scenarios can be represented in a network. Following the investigation of the Anjum case by Crombag and Israëls, our resulting network reached the conclusion that a second scenario, that of Beekman being the killer, is most likely. This was a scenario that was overlooked in the actual trial for this case. The design method is thereby successful in addressing the problem of tunnel vision (1). It should be no surprise that our network reached the same conclusion as Crombag and Israëls, since their work served as a source of information for our model. However, what it clearly showed is that our design method is capable of capturing the narrative approach of Crombag and Israëls in the strong formalization of Bayesian networks. This way, the scenarios are evaluated on a formal account, in which a good (but not true) story will be exposed, addressing (2). Finally, taking the scenarios as a starting point, the method of unfolding helped in keeping an overview of the network and finding all the relevant variables during the building process (3).

A disadvantage of our method lies in the quantification of the entire structure: as in any Bayesian approach, many numbers are difficult to interpret, are unavailable or just disputable. Our method is intended to help build a formal representation of a subjective account of a case, which would help structure the thought process and compare relative probabilities of scenarios. Nonetheless, formalizing subjective probabilities is still not an easy task. A number of elicitation techniques are available, but it remains to be investigated how well these are suited for the specific application to legal cases. Therefore, the elicitation of the required numbers deserves attention in future research.

An opportunity for further developing our design method lies in a more dynamic use of the method. In particular, this could help to actively prevent tunnel vision during the process of investigation. This paper was concerned with the modeling of a collection of relevant scenarios for a case, assuming that a number of fixed scenarios are available. In a more dynamic process, scenarios would be extended or adapted during the process of investigation, simultaneously modeling these changing scenarios in the Bayesian network. Currently, the adaptation of a scenario that was already represented in the Bayesian network is not easily done. A goal for future work is a version of our design method in which various scenarios about what happened can be constructed, extended and adapted in the network during the investigative process.

Acknowledgements This work is part of the project "Designing and Understanding Forensic Bayesian Networks with Arguments and Scenarios" in the Forensic Science programme, financed by the Netherlands Organisation for Scientific Research (NWO). More information about the project: www.ai.rug.nl/˜verheij/nwofs

References

Berger C, Aben D (2010a) Bewijs en overtuiging: Een helder zicht op valkuilen. Expertise en Recht 5/6:159 – 165

Berger C, Aben D (2010b) Bewijs en overtuiging: Rationeel redeneren sinds Aristoteles. Expertise en Recht 2:52 – 56

Berger C, Aben D (2010c) Bewijs en overtuiging: Redeneren in de rechtszaal. Expertise en Recht 3:86 – 90

Bex F (2009) Analysing stories using schemes. In: Kaptein H, Prakken H, Verheij B (eds) Legal Evidence and Proof: Statistics, Stories, Logic, Ashgate Publishing, Aldershot, pp 93 – 116

Bex F (2011) Arguments, Stories and Criminal Evidence, a Formal Hybrid Theory. Springer, Dordrecht

Bex F, van Koppen P, Prakken H, Verheij B (2010) A hybrid formal theory of arguments, stories and criminal evidence. Artificial Intelligence and Law 18:123–152

Conrad JG, Zeleznikow J (2013) The significance of evaluation in AI and law: a case study re-examining ICAIL proceedings. In: Proceedings of the Fourteenth International Conference on Artificial Intelligence and Law, ACM, pp 186–191

Crombag H, Israëls H (2008) Moord in Anjum. Te veel niet gestelde vragen. Boom Juridische uitgevers, Den Haag

Dawid A (2009) Beware of the DAG. In: Journal of Machine Learning Research: Workshop and Conference Proceedings, vol 6, pp 59–86

Fenton N, Neil M (2000) The "jury observation fallacy" and the use of Bayesian networks to present probabilistic legal arguments. Mathematics Today 36(6):180 – 187

Fenton N, Neil M (2012) On limiting the use of Bayes in presenting forensic evidence. Available online: `http://www.eecs.qmul.ac.uk/~norman/papers/likelihood_ratio.pdf`

Fenton N, Neil M, Lagnado D (2011) Modelling mutually exclusive causes in Bayesian networks. Available online: `http://www.eecs.qmul.ac.uk/~norman/papers/mutual_IEEE_format_version.pdf`

Fenton N, Neil M, Lagnado D (2013) A general structure for legal arguments using Bayesian networks. Cognitive Science 37:61 – 102

van Gosliga S, van de Voorde I (2008) Hypothesis management framework: A flexible design pattern for belief networks in decision support systems.

In: 6th Bayesian Modelling Applications Workshop on UAI 2008, Helsinki, Finland
Handfield T (2012) A Philosophical Guide to Chance: Physical Probability. Cambridge University Press
Hepler A, Dawid A, Leucari V (2004) Object-oriented graphical representations of complex patterns of evidence. Law, Probability & Risk 6:275–293
Jensen F, Nielsen T (2007) Bayesian Networks and Decision Graphs. Springer, New York
Kaptein H, Prakken H, Verheij B (2009) Legal Evidence and Proof: Statistics, Stories, Logic. Ashgate Publishing Company, Aldershot
Keppens J (2011) On extracting arguments from Bayesian network representations of evidential reasoning. In: Ashley K, van Engers T (eds) The 13th International Conference on Artificial Intelligence and Law, ACM, New York, pp 141 – 150
Keppens J, Schafer B (2006) Knowledge based crime scenario modelling. Expert Systems with Applications 30(2):203–222
Lagnado D, Fenton N, Neil M (2013) Legal idioms: a framework for evidential reasoning. Argument and Computation 4(1):46–63
Laskey K, Mahoney S (1997) Network fragments: Representing knowledge for constructing probabilistic models. In: Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann Publishers Inc., pp 334–341
Pearl J (1988) Embracing causality in default reasoning. Artificial Intelligence 35:259–271
Pennington N, Hastie R (1992) Explaining the evidence: Tests of the story model for juror decision making. Journal of Personality and Social Psychology 62(2):189 – 206
Pennington N, Hastie R (1993) Inside the Juror, Cambridge University Press, Cambridge, chap The Story Model for Juror Decision Making, pp 192 – 221
Poot CD, Bokhorst R, van Koppen P, Mulder E (2004) Recherche Portret: Over Dilemma's in de Opsporing. Kluwer, Alphen aan den Rijn
Prakken H (2010) An abstract framework for argumentation with structured arguments. Argument and Computation 1:93 – 124
Renooij S (2001) Probability elicitation for belief networks: Issues to consider. The Knowledge Engineering Review 16(3):255 – 269
Rumelhart D (1975) Notes on a schema for stories. In: Bobrow D, Collins A (eds) Representation and Understanding: Studies in Cognitive Science, Academic Press, New York, New York
Schank R, Abelson R (1977) Scripts, Plans, Goals and Understanding, An Inquiry into Human Knowledge Structures. Lawrence Erlbaum, Hillsdale
Sileno G, Boer A, van Engers T (2012) Analysis of legal narratives: a conceptual framework. In: Legal Knowledge and Information Systems: JURIX 2012: the 25th Annual Conference, IOS Press, Amsterdam, pp 143–146
Taroni F, Aitken C, Garbolino P, Biedermann A (2006) Bayesian Networks and Probabilistic Inference in Forensic Science. Wiley, Chichester

Timmer S, Meyer JJ, Prakken H, Renooij S, Verheij B (2013) Inference and attack in Bayesian networks. In: Hindriks K, de Weerdt M, van Riemsdijk B, Warnier M (eds) Proceedings of the 25th Benelux Conference on Artificial Intelligence, pp 199 – 206

Vlek C, Prakken H, Renooij S, Verheij B (2013a) Modeling crime scenarios in a Bayesian network. In: Proceedings of the 14th International Conference on Artificial Intelligence and Law, ACM Press, New York, pp 150 – 159

Vlek C, Prakken H, Renooij S, Verheij B (2013b) Representing and evaluating legal narratives with subscenarios in a Bayesian network. In: Finlayson M, Fisseni B, Löwe B, Meister J (eds) 2013 Workshop on Computational Models of Narrative, Schloss Dagstuhl, Saarbrücken/Wadern, Germany, pp 315–332, DOI 10.4230/OASIcs.CMN.2013.i

Vlek C, Prakken H, Renooij S, Verheij B (2013c) Unfolding crime scenarios with variations: A method for building bayesian networks for legal narratives. In: Ashley K (ed) Legal Knowledge and Information Systems: JURIX 2013: The Twenty-Sixth Annual Conference, IOS Press, pp 145 – 154

Wagenaar W, van Koppen P, Crombag H (1993) Anchored Narratives: the Psychology of Criminal Evidence. Harvester Wheatsheaf, Hemel Hempstead

8 Appendix: additional figures

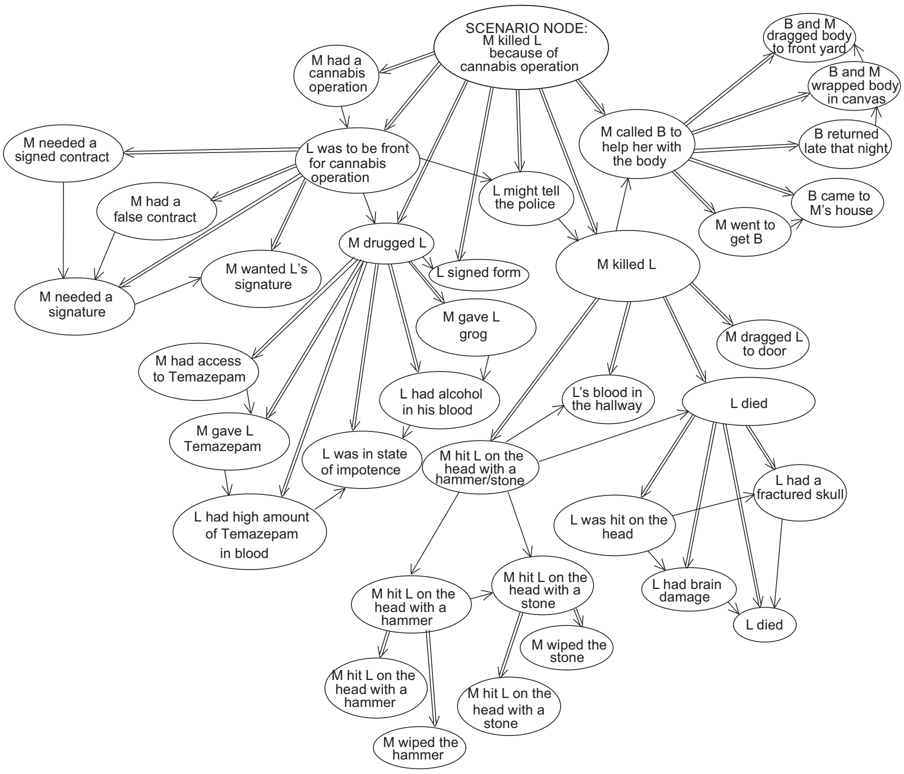

Fig. 18: The first scenario

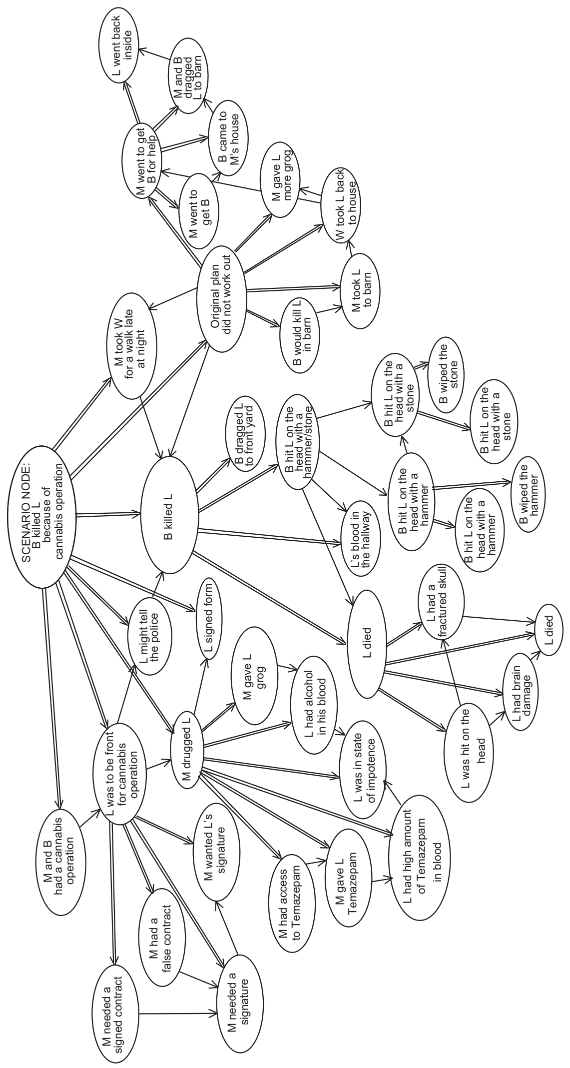

Fig. 19: The second scenario

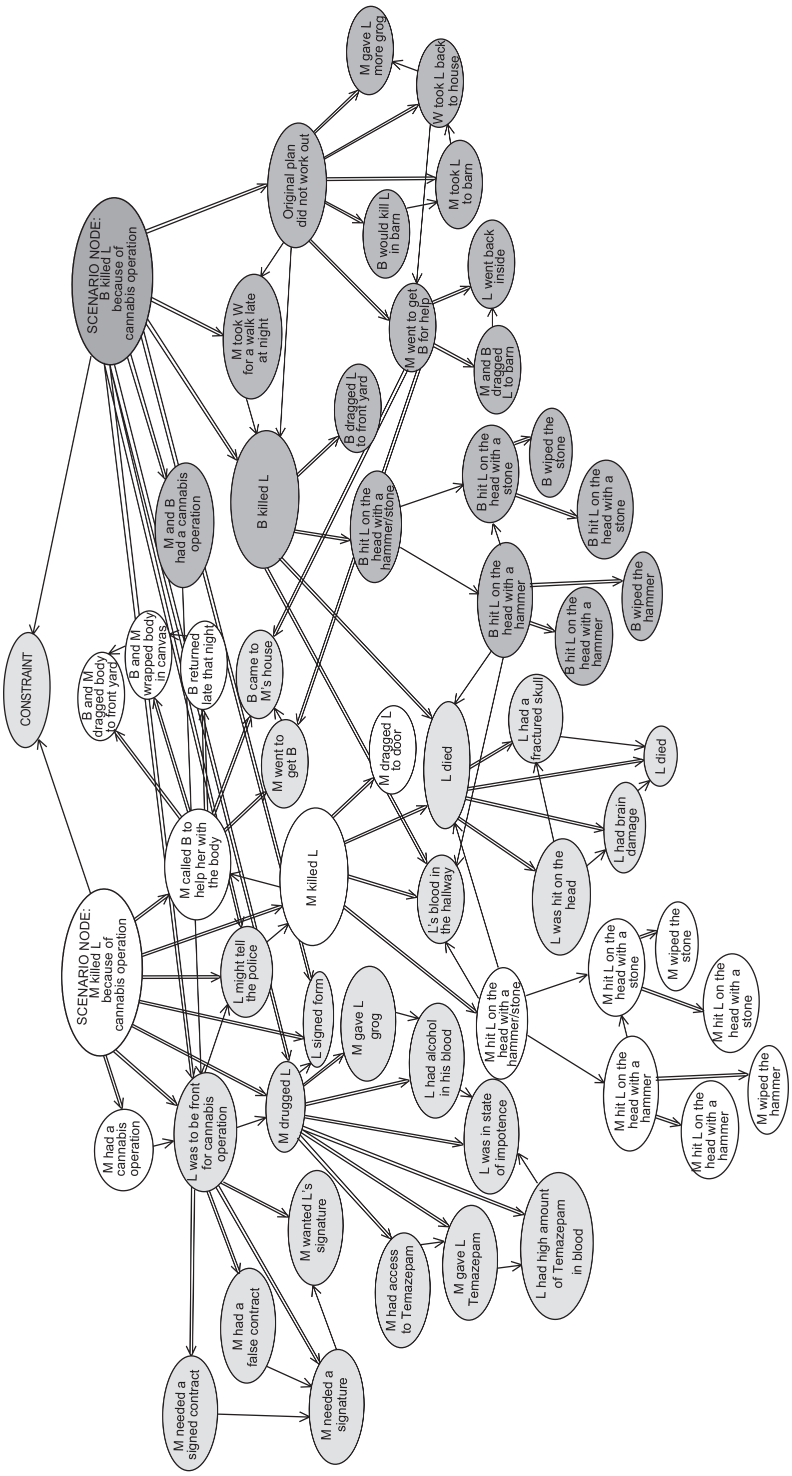

Fig. 20: The merged scenarios. The nodes of the first scenario only are white, nodes of the second scenario only are dark grey, and nodes in both scenarios are light grey.

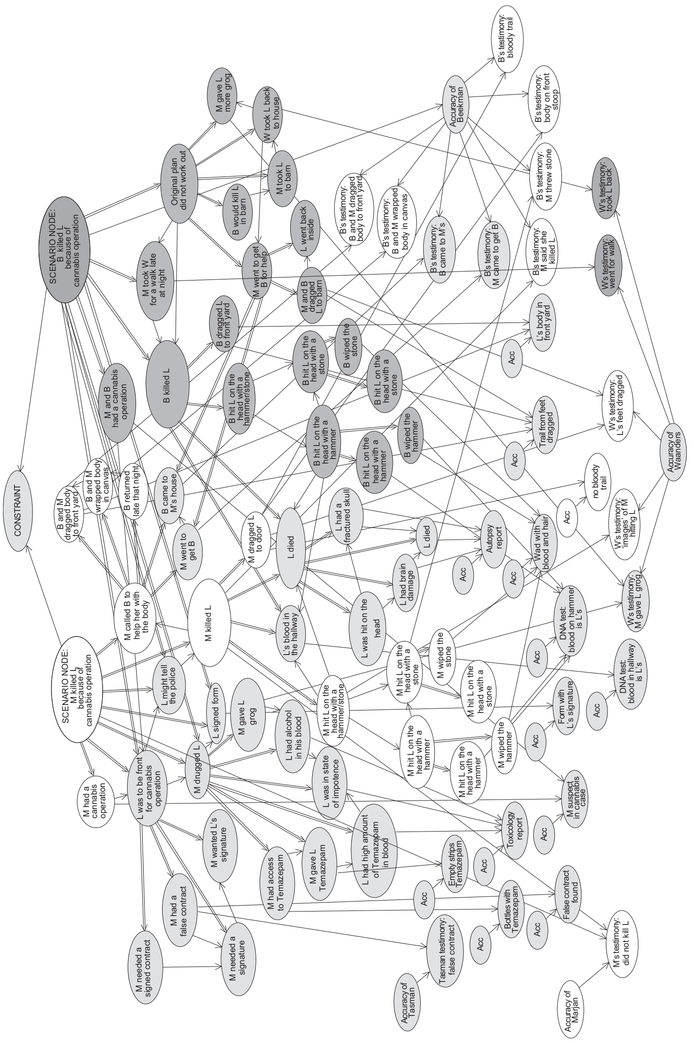

Fig. 21: The full structure with the evidence. Nodes of the first scenario only are white, nodes of the second scenario are dark grey, and nodes in both scenarios are light grey.