Representing the quality of crime scenarios in a Bayesian network

Charlotte S. VLEK [a,1], Henry PRAKKEN [b,c], Silja RENOOIJ [b] and Bart VERHEIJ [a]
[a] *Institute of Artificial Intelligence, University of Groningen*
[b] *Department of Information and Computing Sciences, Utrecht University*
[c] *Faculty of Law, University of Groningen*

Abstract
Bayesian networks have gained popularity as a probabilistic tool for reasoning with legal evidence. However, two common difficulties are (1) the construction and (2) the understanding of a network. In previous work, we proposed to use narrative tools and in particular scenario schemes to assist the construction and the understanding of Bayesian networks for legal cases. We proposed a construction method and a reporting format for explaining or understanding the network. The quality of a scenario, which plays an important role in the narrative approach to evidential reasoning, was not yet included in this method. In this paper, we provide a discussion of what constitutes the quality of a scenario, in terms of the narrative concepts of completeness, consistency and plausibility. We propose a probabilistic interpretation of these concepts, and show how they can be incorporated in our previously proposed method. We also illustrate with an example how these concepts concerning scenario quality can be used to explain or understand a Bayesian network.

Keywords. Reasoning about legal evidence, Bayesian networks, narrative

Introduction

When reasoning about evidence in a legal case, Bayesian networks can provide a good tool for a probabilistic analysis. Currently, Bayesian networks are typically used for analysing parts of a case that are specifically concerned with probabilistic evidence (see [1]). Little work has been done on analysing an entire legal case in a Bayesian network, though a notable exception is [2]. A Bayesian network is potentially a suitable candidate for analysing a case as a whole, since the network deals well with the various dependencies and independencies a case might contain. However, in practice, analysing an entire case in a Bayesian network is not a simple task, mainly due to two issues: Bayesian networks are often hard to construct and hard to understand. In our work, we aim to address these issues by combining Bayesian networks with a narrative approach.

In previous work, we proposed the use of scenarios to guide the construction of a Bayesian network from narrative idioms as basic building blocks [3]. In [4] we recognised that scenario schemes can not only provide further structure to the construction of a Bayesian network, but can also aid the understanding of the model. This resulted in a

[1]Corresponding Author: Charlotte S. Vlek, University of Groningen, E-mail: c.s.vlek@rug.nl.

method for constructing a Bayesian network capturing the various scenarios of a case, and a method for reporting the scenarios from the network and their relations to the evidence. What was lacking, however, was a treatment of the quality of scenarios in the modelling and in the understanding of the network.

In the literature on scenario-based reasoning, there is no clear definition of the quality of a scenario. Nonetheless, various authors emphasize the role of the quality of a scenario in the final acceptability of a scenario [5,6,7]. Although the quality of a scenario as a concept remains somewhat vague, separate factors that make a good scenario are discussed in the literature. In this paper we explore three such factors as proposed in [6] and later also adopted in [7]: completeness, consistency and plausibility. We discuss each of these factors, and how they can be interpreted in a probabilistic setting. This way, the three factors can be included in the construction method from [4] to enable the representation of scenario quality in a Bayesian network.

With a fictional example case, we then illustrate how these three factors can enhance the understanding of the network. In particular, by understanding which elements of a scenario are plausible or implausible, it becomes apparent why certain evidence has a crucial impact on the probability of the scenario (since it supports an implausible element), or which evidence is crucially lacking (creating so-called evidential gaps).

In sum, the contributions of this paper are as follows: we discuss the factors that determine a scenario's quality, and how they can be represented using the method from [4]. We then illustrate how these factors can be used for a better understanding of the resulting Bayesian network for a case.

1. Preliminaries

In this section, some background is discussed on Bayesian networks (Section 1.1) and on previous work concerning the construction and understanding of Bayesian networks (Section 1.2).

1.1. Bayesian networks

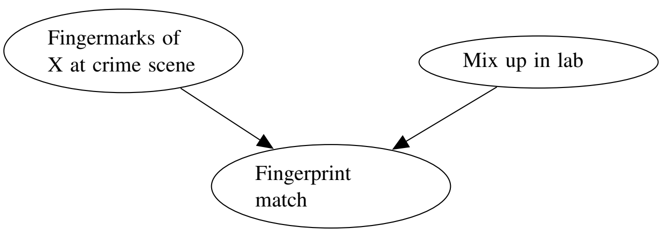

Figure 1. An example of a Bayesian network graph

A Bayesian network is a compact representation of a joint probability distribution (JPD) [8]. It consists of a directed acyclic graph, such as the one shown in Figure 1, and probability tables for each node in the graph. Each node can have certain values (e.g. true/false, or more than two values). The probability table of a node V with values v_1, v_2, v_3 expresses the probability of V taking value v_i (for each value v_i) conditioned on any configuration of values of the set of parents of V. The probability table for `Fingerprint match` would thus contain the probabilities of this node taking value true/false conditioned on any combination of values of `Fingermarks of X at crime scene` and `Mix up in lab`. When a node has no parents, the probability table specifies the prior probability for each value (e.g. the prior probability `Mix up in lab = T`).

From the arrows in the graph, information about (in)dependences between variables can be read. Arrows are often drawn such that they represent causality, although technically the arrows are not necessarily causal relations [9]. Once a Bayesian network has been fully specified, any prior or posterior probability of interest can be calculated from the network. In legal applications, what is typically of interest is the probability of a certain hypothesis $H = \mathtt{T}$ given the evidence $E_1 = e_1, ...E_n = e_n$. In Figure 1 the hypothesis of interest may be `Fingermarks of X at crime scene = T`, with evidence that a fingerprint match was found (`Fingerprint match = T`).

1.2. Constructing and understanding Bayesian networks with scenario schemes

In this section we briefly discuss our previous work [4], in which we proposed to use scenario schemes to assist the construction and understanding of a Bayesian network.

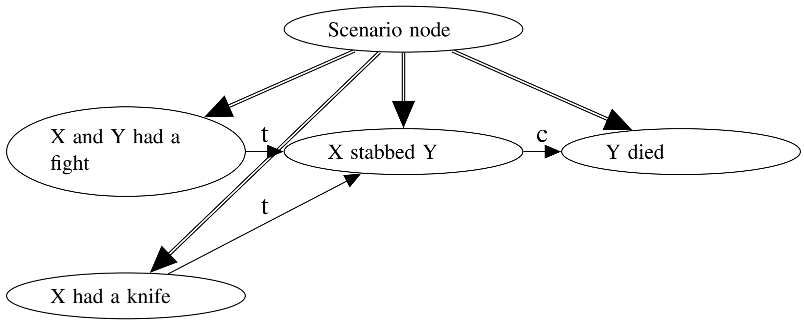

Figure 2. A scenario scheme for a fight resulting in stabbing

Scenario scheme idioms A scenario scheme gives the abstract structure for a typical scenario [7], such as, for example, a typical burglary or murder. Based on this concept we proposed the use of scenario scheme idioms, which are Bayesian network structures of nodes and arrows, representing a scenario scheme. These idioms can be used as a building block for constructing a Bayesian network, similar to the idioms from [10] and [3]. An example of a scenario scheme idiom is shown in Figure 2, about two people getting into a fight, resulting in the stabbing and death of one of them.

The graphical structure of a scenario scheme idiom builds on the ideas from [3], in which narrative idioms were proposed to capture the coherent structure of a scenario. The structure of a scenario scheme idiom is always such that the scenario as a whole is represented with a so-called *scenario node*, with the elements of the scenario scheme (X stabbed Y, etcetera) as child nodes of the scenario node. The connection between the scenario node and each element node ensures an influence between these elements, thereby capturing the coherence of the scenario (see [3] for details on this construction). Ultimately, the goal is for a modeller to have a database of scenario scheme idioms available when building a Bayesian network based on a collection of scenarios. However, even when such a database is not available and the modeller has to design the idioms herself, the sheer task of explicating which scenario scheme underlies a scenario already helps the modelling process.

Retrieving a scenario in text form from the network Due to the structure of a scenario scheme idiom, the resulting network has certain properties. Firstly, elements of the scenario (propositions) correspond to variables in the network when assigned value true. Secondly, these variables are formulated such that an arrow between nodes signifies a positive influence between propositions. Thirdly, the arrows are labelled to indicate causal connections ('c') and temporal connections ('t'). These properties enable us to retrieve from the network a scenario in text form. A method was proposed in [4] to do this. This method used the direction of the arrows to decide on the order of propositions and the labels on the arrows to add connectives to the text.

Reporting the scenarios and the evidence After retrieving the scenarios in text form from the network, they are used in a report about the network. This report contains a list of the evidence, and how strongly it supports or attacks each scenario. The strength of support is reported verbally by converting a numerical measure with a verbal scale. A combined strength of the evidence is also reported, to incorporate possible dependencies between the evidence. In addition, some narrative concepts are used to provide more insight. This includes pointing out *distinguishing evidence* (pieces of evidence that support one scenario more than another) and *scenario consequences* (evidence that would be expected as a consequence of a scenario).

2. Representing the quality of a scenario

In this section we show how the quality of a scenario can be represented with the method reviewed in Section 1.2. To this end, we first discuss three concepts concerning a scenario's quality, taken from the literature on scenario-based reasoning with legal evidence (Section 2.1). Then we discuss how to represent these concepts (Section 2.2).

2.1. Narrative concepts on the quality of a scenario

According to the Anchored Narratives theory [5], a good scenario should be unambiguous, contain no contradictions and have a central action that is logical in the setting of the scenario. Similar properties were described by Pennington and Hastie [6] and Bex [7]. In what follows, we will use the (shared) terminology from [6] and [7], in which the quality (or *coherence*) of a scenario depends on three factors: completeness, consistency and plausibility.

Completeness A scenario is complete when it 'has all its parts' [6]. For example, Pennington and Hastie describe a typical scenario about an intentional action, which should contain initiating states and events, goals, actions and consequences. Bex [7] formalised this using scenario schemes: a scenario is complete when it corresponds to a plausible scenario scheme. To correspond to a scheme, the scenario should *complete* the scheme (each element of the scheme has a corresponding element in the scenario) as well as *fit* the scheme (each element of the scenario has a corresponding element in the scheme). According to Bex, a scenario should thus 'have all its parts', yet no 'extra parts'.

Consistency For a scenario to be consistent it should not contain internal contradictions with the evidence or with other parts of the scenario (see [5], [6] and [7]). Such internal contradictions can be fairly obvious (e.g. if someone is at home, he cannot be at the supermarket as well), or require some inference steps (if someone drove somewhere with a car, this car should be at that location too). When a scenario is inconsistent, it cannot be a good scenario and should not be considered as a possible alternative.

Plausibility A scenario is ‘plausible to the extent that it corresponds to the decision maker’s knowledge about what typically happens in the world and does not contradict that knowledge’ [6]. A scenario can thus be plausible to a certain degree. In Bex [7], the plausibility of a scenario depends on its elements and connections (or generalizations), and how well they are supported by arguments based on common sense knowledge.

As remarked in [5] but also by others, various elements of a scenario require various levels of support (or anchoring, in terms of [5]). With proper support, an implausible scenario can become quite believable. The latter leads to an interesting effect of (im)plausibility which we will try to capture in our probabilistic models: support for an implausible element of a scenario can have a strong effect on the overall probability of that scenario, which may be much stronger than the effect of support for a plausible element. For example, consider the following scenario: ‘Jane and Mark had a fight. Jane had a knife, and she threatened Mark with it. Mark hit Jane, who then dropped the knife. Mark fell on the knife and died.’ This is at first sight an implausible scenario, but if the event that Mark fell on the knife is somehow supported with convincing evidence, the entire scenario can go from very improbable to quite probable. On the other hand, evidence for Jane having a knife has much less effect, since this was already more credible.

2.2. *Representing the quality of a scenario in a Bayesian network*

In this section, the properties completeness, consistency and plausibility as discussed in the previous subsection are investigated in the context of Bayesian networks. In particular, we investigate how these properties can be represented with the method from [4].

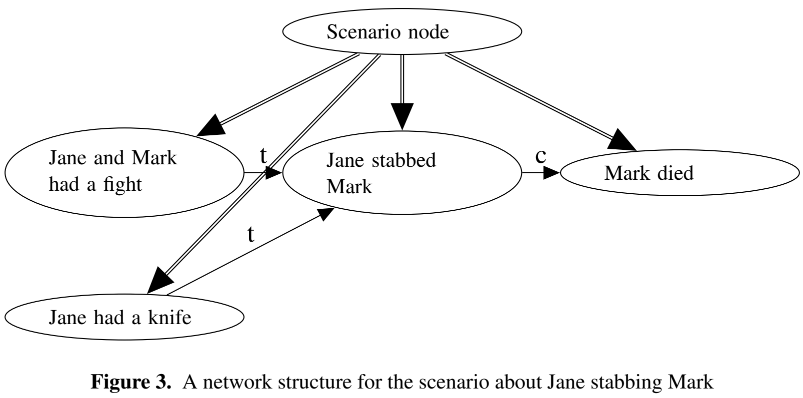

Figure 3. A network structure for the scenario about Jane stabbing Mark

Representing completeness In the method from [4], when an incomplete scenario is encountered during the modelling process, the modeller uses an appropriate scenario scheme to add the parts of the scenario that are missing, thereby filling the gaps to make it a complete scenario. Using Bex's formalization of completeness, extending a scenario such that it fits a scenario scheme means that it is by definition now a complete scenario.

As an example of an incomplete scenario, consider the following: 'Jane stabbed Mark, Mark died'. This scenario could fit the scenario scheme from Figure 2, but it lacks a motive for the stabbing. Completing the scenario with this scheme would lead to a network structure as shown in Figure 3. However, the same incomplete scenario could also fit a scenario scheme in which Jane stabbing Mark was self defence, in which case it lacks circumstances to explain why Jane needed to defend herself. When the modeller finds both extensions (or even more than two) to be relevant for the case, she can choose to model them as alternative scenarios. Note that the modeller does not need to be certain that the circumstances (e.g. Jane having to defend herself) took place, since it is simply modelled as a variable in the network that can be true or false.

Representing consistency A consistent scenario can be modelled using a scenario scheme idiom. If a scenario is inconsistent, this can be represented with a constraint node to model that two events cannot occur together (see [11] for the constraint construction).

Suppose a scenario needs to be modelled that contains the events 'Jane took off in a car' and 'Jane took off on foot'. This scenario contains an internal inconsistency and is therefore not an acceptable alternative. A constraint node is added with arrows from the inconsistent events to the constraint node. Due to the properties of the scenario scheme idiom (see [4]) the scenario as a whole is now no longer an alternative: the constraint node construction leads to a probability of 0 for the scenario that contains the inconsistency.

Representing plausibility As discussed in Section 2.1, the plausibility of a scenario depends on how well it is supported by our common sense knowledge about the world. In a Bayesian network, the probability tables represent our common sense knowledge about the world, without taking any evidence into account. The prior probability of the scenario node being true can thus be viewed as expressing the plausibility of that scenario. However, plausibility can provide more structure to the construction process by not only looking at a scenario's *global* plausibility (in terms of the probability of the scenario node being true), but also at the plausibility of specific elements or connections within the scenario, similar to [7]. When an element or connection is implausible, this will be reflected in the network with a low probability for that element or for an element conditioned on another element (for a connection).

Consider the implausible event 'Mark fell on a knife'. Its implausibility is represented in a Bayesian network with a low probability $\Pr(\texttt{Mark fell on a knife = T})$. However, this number is not present in a probability table, so instead we need to set the probability $\Pr(\texttt{Mark fell on a knife = T}|\texttt{Scenario node = F})$ to an appropriate value. This can be calculated based on the value that we want $\Pr(\texttt{Mark fell on a knife = T})$ to take, as follows (using that the probabilities conditioned on `Scenario node = T` are set to 1 by the method from [4]):

$$\Pr(\texttt{Mark fell on a knife = T}|\texttt{Scenario node = F})$$
$$= \frac{\Pr(\texttt{Mark fell on a knife = T}) - \Pr(\texttt{Scenario node = T})}{\Pr(\texttt{Scenario node = F})}$$

Consider the connection between ‘Mark fell on a knife’ and ‘Mark died’, which is considered implausible because it would be very unlikely that the knife was positioned in such a way that falling on it would result in a fatal injury. Such an implausible connection dictates that the probability for Pr(`Mark died = T`|`Mark fell on a knife = T`) is low. Similar to the previous example, this needs to be established by setting an appropriate value for Pr(`Mark died = T`|`Mark fell on a knife = T`, `Scenario node = F`) in the probability table. Again, this can be calculated from Pr(`Mark died = T`|`Mark fell on a knife = T`), assuming that the probabilities for `Mark fell on a knife` are already known (abbreviating the scenario node to `ScN` and `Mark fell on a knife` to `Mark fell...`):

$$\Pr(\texttt{Mark died = T}|\texttt{Mark fell on a knife = T}, \texttt{ScN = F})$$
$$= \frac{\Pr(\texttt{Mark died = T}|\texttt{Mark fell... = T}) \cdot \Pr(\texttt{Mark fell...= T}) - \Pr(\texttt{ScN = T})}{\Pr(\texttt{Mark fell...= T}|\texttt{ScN = F}) \cdot \Pr(\texttt{ScN = F})}$$

As described in Section 2.1, an interesting effect of (im)plausibility is that evidential support for an implausible element of a scenario can have a stronger effect on the scenario as a whole than support for a plausible element. This property of implausibility is indeed captured by representing implausible events and connections as described above: by giving them a low probability. To see this, consider an implausible event A (Mark falling on a knife), with a low probability Pr(A = `T`). As soon as we learn (via some evidence) that the event A is likely true, it follows that the scenario node is probably *not* false (also recall that the probability conditioned on scenario node = true is 1 for any element of the scenario). However, for an event B (Jane having a knife) that is quite plausible, the probability Pr(B = `T`) is higher, which makes it less informative about the value of the scenario node: the scenario node might very well still be false while B is true.

3. Understanding the quality of a scenario: an example

In this section we illustrate how the ideas from Section 2 can be used to explain or understand the quality of a scenario in a Bayesian network. This serves as an addition to our report as proposed in [4], in which the posterior probability of each scenario, as well as the relations to the evidence were reported. As an example, we model the fictional case of Jane and Mark that has already served as an illustration in the previous sections.

The following scenario was modelled in Section 2: ‘Jane and Mark had a fight, Jane then stabbed Mark and Mark died’ (Figure 3). Another scenario was also briefly discussed: ‘Jane and Mark had a fight. Jane had a knife, and she threatened Mark with it. Mark hit Jane, who then dropped the knife. Mark fell on the knife and died.’ This can be modelled with a scenario scheme idiom about Mark dying by accident, and a scenario scheme idiom to further specify the accident with a subscenario. A model of the second scenario as shown in Figure 4. Combining Figures 3 and 4 (using the idioms from [3]) leads to a full network for the case. As discussed in Section 2, the (im)plausibility of elements of the scenario can serve as a guideline for eliciting some of the required probabilities. This leads to low probabilities (say, below 0.01) for Pr(`Mark fell on a knife = T`) and Pr(`Mark died = T`|`Mark fell on a knife = T`). Other probabilities can be subjectively established using elicitation techniques (see e.g. [12]).

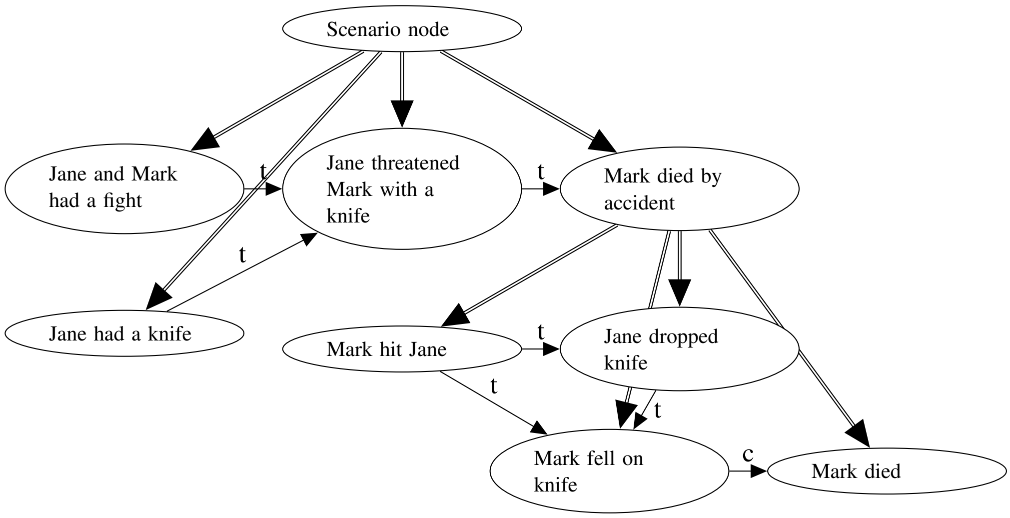

Figure 4. A Bayesian network structure for the second scenario

Understanding completeness During the construction of the network, any incomplete scenario that is encountered will be extended such that it fits a scenario scheme, thereby making it complete. As a result, the scenarios that are modelled in a Bayesian network must be complete. This is also the case in our example.

Understanding consistency In this case, no inconsistency was encountered in the scenarios. However, it is not hard to imagine a third, inconsistent scenario (such as: Jane did not have a knife, but she did stab Mark), which is then represented with a constraint node. This would result in a posterior probability of 0 for the scenario node to be true. By recognizing the constraint node construction and explaining that the two elements of the scenario cannot occur simultaneously, the probability of 0 can be understood.

Understanding plausibility As described in Sections 2.1 and 2.2, an effect of plausibility can be that support for an implausible element of the scenario has a stronger effect on the posterior probability of the scenario as a whole than support for a plausible event. Understanding that an implausible element or connection receives evidential support is thus informative to understand the outcome of the network. Moreover, unsupported implausible elements or connections are also of interest: these form the so-called *evidential gaps* of a scenario, and they reveal which crucial points in the scenario are unsupported.

An implausible event A can be identified in a network by its low probability $\Pr(A = \texttt{T})$. Similarly, an implausible connection between events A and B is identified by a low probability for $\Pr(B = \texttt{T}|A = \texttt{T})$. In what follows, we assume that implausibility not only leads to low probabilities, but also that a low probability for a variable to have value true is always indicative of an implausible event or connection. From the probabilities one can now easily find the required probabilities, and determine whether they are below some threshold (say, 0.01), making it implausible.

In our example case, the event 'Mark fell on a knife' was modelled to have a low probability, since it is an implausible event. As long as there is no evidence for this event, it is an evidential gap: it is a crucial element that remains unsupported by evidence. This crucial lack of evidence shows why the second scenario has a low posterior probability.

As soon as there is convincing evidence for the event that Mark fell on a knife (e.g. a policeman saw him fall), this is no longer an evidential gap, but an implausible element with evidential support. By understanding that such an implausible element is now supported, a judge or jury can see why the scenario that was implausible without evidence, now became (much) more probable with the evidence.

4. Discussion and conclusion

In this paper, an analysis of scenario quality was presented, to enable the representation and understanding of scenario quality in our previously proposed method from [4]. We provided a probabilistic interpretation of the three factors that determine the quality of a scenario: completeness, consistency and plausibility. We showed how these three factors can be represented in a Bayesian network with the method from [4], and we illustrated how these notions can be applied to the understanding of the Bayesian network.

Completeness can be represented in a Bayesian network by always using a scenario scheme idiom that fits the scenario. This way, only complete scenarios are modelled. Our interpretation of completeness remains close to how it was originally formulated by Pennington and Hastie [6], and later interpreted by Bex [7], using scenario schemes. Pennington and Hastie speak of a scheme to model intentional actions, while Bex allows for various context-dependent schemes, which is something we adopt as well.

When two elements are inconsistent, this can be represented with a constraint node. As a result, the probability of the scenario as a whole becomes 0, which means that an inconsistent scenario is not taken into account as a possible alternative. In this interpretation of consistency we follow the ideas by Pennington and Hastie and Bex, that a scenario is either consistent or inconsistent and when it is inconsistent it should not be taken into account as an alternative. However, we remark that a probabilistic interpretation lends itself to a notion of degrees of consistency, rather than the current black-and-white interpretation in which a scenario is either consistent or inconsistent. For example, a so-called conflict measure [8] can express how consistent two pieces of evidence are together. It might be possible to extend this notion to report not only on conflicting evidence, but also on conflicting elements in a scenario. This may be of interest for future work, in which we intend to develop methods for reporting on scenario quality.

Finally, we proposed to represent plausibility in the probability tables of the network. The global plausibility of a scenario as a whole can be represented in the prior probability of the scenario node, while specific implausible events or connections within the scenario are represented with a low probability for that element or connection. This interpretation is somewhat farther removed from the ideas in [6] and [7], since we explicitly employ probabilities. While plausibility remains somewhat informal in [6] and [7], the authors seem to be in agreement that plausibility has to do with our background knowledge about the world. In [7], this leads to an interpretation of plausibility in which arguments based on common sense knowledge can support or attack elements of the scenario. In our interpretation, common sense knowledge is not as explicitly present as in Bex's work, but arguably, the probabilities in the probability tables of a Bayesian network represent how we commonly view the world, thus representing common sense knowledge.

Thinking about the quality of a scenario can assist the construction of a network as well as the understanding. By making explicit whether a scenario is complete, the modeller is encouraged to choose the correct scenario scheme to which a scenario corresponds. Furthermore, thinking about the plausibility of elements provides some heuristics for eliciting the numbers needed for the probability tables of a Bayesian network.

As was illustrated with the example, understanding implausible events or connections can also help to understand the results of a Bayesian network. This is because such implausible elements can play an important role in deciding whether a scenario is acceptable: when an implausible element is supported by evidence, this can have great impact on the posterior probability of a scenario, while an implausible element that remains unsupported can be pointed out as an evidential gap. In the future, we aim to work out these ideas further and develop a reporting format for reporting on scenario quality.

Our work on narrative and Bayesian networks was prompted by a more fundamental desire to understand the connections between three prominent formalisms for reasoning with legal evidence: the probabilistic approach, the narrative approach and the argumentative approach (see also [13]). Related work on formalising scenario-based reasoning was done by Bex [7], who developed a formal model for reasoning with scenarios and arguments. The connection between arguments and Bayesian networks is investigated by Timmer et. al. [14]. The contribution of this paper is to formalise the quality of a scenario in a probabilistic setting.

References

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

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

[3] C. Vlek, H. Prakken, S. Renooij, and B. Verheij, “Building Bayesian networks for legal evidence with narratives: a case study evaluation,” *Artifical Intelligence and Law*, vol. 22, no. 4, pp. 375 – 421, 2014.

[4] C. Vlek, H. Prakken, S. Renooij, and B. Verheij, “Constructing and understanding Bayesian networks for legal evidence with scenario schemes,” in *Proceedings of the 15th International Conference on Artificial Intelligence and Law*, (New York), pp. 128– 137, ACM Press, 2015.

[5] W. Wagenaar, P. van Koppen, and H. Crombag, *Anchored Narratives: the Psychology of Criminal Evidence*. Hemel Hempstead: Harvester Wheatsheaf, 1993.

[6] N. Pennington and R. Hastie, *Inside the Juror*, ch. The Story Model for Juror Decision Making, pp. 192 – 221. Cambridge: Cambridge University Press, 1993.

[7] F. Bex, *Arguments, Stories and Criminal Evidence, a Formal Hybrid Theory*. Dordrecht: Springer, 2011.

[8] F. Jensen and T. Nielsen, *Bayesian Networks and Decision Graphs*. New York: Springer, 2007.

[9] A. Dawid, “Beware of the DAG,” in *Journal of Machine Learning Research: Workshop and Conference Proceedings*, vol. 6, pp. 59–86, 2009.

[10] N. Fenton, M. Neil, and D. Lagnado, “A general structure for legal arguments using Bayesian networks,” *Cognitive Science*, vol. 37, pp. 61 – 102, 2013.

[11] N. Fenton, M. Neil, and D. Lagnado, “Modelling mutually exclusive causes in Bayesian networks,” 2011. Available online: http://www.eecs.qmul.ac.uk/∼norman/papers/mutual_IEEE_format_version.pdf.

[12] S. Renooij, “Probability elicitation for belief networks: Issues to consider,” *The Knowledge Engineering Review*, vol. 16, no. 3, pp. 255 – 269, 2001.

[13] B. Verheij, F. Bex, S. Timmer, C. Vlek, J.-J. Meyer, S. Renooij, and H. Prakken, “Arguments, scenarios and probabilities: Connections between three normative frameworks for evidential reasoning,” *Law, Probability & Risk*, 2015. To appear.

[14] S. Timmer, J.-J. Meyer, H. Prakken, S. Renooij, and B. Verheij, “A structure-guided approach to capturing Bayesian reasoning about legal evidence in argumentation,” in *Proceedings of the 15th International Conference on Artificial Intelligence and Law*, (New York), pp. 109– 118, ACM Press, 2015.