Supporting Discussions About Forensic Bayesian Networks Using Argumentation

Remi Wieten
Department of Information and Computing Sciences,
Utrecht University, The Netherlands
g.m.wieten@uu.nl

Floris Bex
Department of Information and Computing Sciences,
Utrecht University, The Netherlands
Institute for Law, Technology and Society,
Tilburg University, The Netherlands
f.j.bex@uu.nl

Henry Prakken
Department of Information and Computing Sciences,
Utrecht University, The Netherlands
Faculty of Law, University of Groningen, The Netherlands
h.prakken@uu.nl

Silja Renooij
Department of Information and Computing Sciences,
Utrecht University, The Netherlands
s.renooij@uu.nl

ABSTRACT

Bayesian networks (BNs) are powerful tools that are increasingly being used by forensic and legal experts to reason about the uncertain conclusions that can be inferred from the evidence in a case. Although in BN construction it is good practice to document the model itself, the importance of documenting design decisions has received little attention. Such decisions, including the (possibly conflicting) reasons behind them, are important for legal experts to understand and accept probabilistic models of cases. Moreover, when disagreements arise between domain experts involved in the construction of BNs, there are no systematic means to resolve such disagreements. Therefore, we propose an approach that allows domain experts to explicitly express and capture their reasons pro and con modelling decisions using argumentation, and that resolves their disagreements as much as possible. Our approach is based on a case study, in which the argumentation structure of an actual disagreement between two forensic BN experts is analysed.

CCS CONCEPTS

• **Applied computing** → **Law**; • **Computing methodologies** → Knowledge representation and reasoning; • **Mathematics of computing** → Probability and statistics.

KEYWORDS

Bayesian networks, argumentation support, legal reasoning

1 INTRODUCTION

Bayesian networks (BNs) are powerful tools that are well-suited for reasoning about the uncertain consequences that can be inferred from the evidence in a legal case [1]. A BN compactly represents a joint probability distribution by means of a graph-structure and tables with probability parameters. In recent years, forensic and legal experts have increasingly developed and used BNs for the interpretation of different types of forensic trace evidence [2], such as glass fragments, finger marks and DNA evidence, as well as entire legal cases [3]. Although in BN construction it is good practice to document the model itself, the importance of documenting design decisions has received little attention. Such decisions, including the (possibly conflicting) reasons behind them, are important for legal experts to understand and accept probabilistic models of cases. Moreover, when disagreements arise between domain experts involved in the construction of BNs, there are no systematic means to resolve such disagreements.

An example of a tool that does support domain experts in documenting their BN modelling decisions is that of Yet and colleagues [4]. Their tool allows BN developers to document the (clinical) evidence, including conflicts, underlying the constructed BN in a queryable OWL ontology, which users can examine through an automatically generated free text web page. The tool, however, does not support (domain) experts in recording conflicting reasons, let alone that it provides experts the ability to resolve them. Since the legal domain is inherently adversarial, we prefer to use an approach for capturing and resolving such conflicts based on argumentation [5]. Keppens [6] recently proposed such an argumentation-based approach to criticise and resolve discussions about a probability distribution. Keppens' approach can be used when discussions about a BN only concern the parameterisation of a fully specified BN. However, disagreements also arise regarding partially specified BNs, and may also concern a BN's graph-structure.

Accordingly, in this paper we propose an approach that helps to capture and resolve disagreements among experts concerning any

BN element. To this end, we allow experts to explicitly express their reasons pro and con modelling decisions regarding the structure and parameterisation of a (fully or partially specified) BN using argumentation. Disagreements are then resolved as much as possible by utilising preferences that are specified over the arguments by the experts. Our approach is based on a case study, in which we analyse the argumentation structure of an actual disagreement between two experts at the Netherlands Forensic Institute about a forensic BN for the interpretation of finger marks. In a recent case study, Prakken [7] analysed the argumentation structure of actual court discussions regarding Bayesian analyses of criminal cases. However, he mainly concerned himself with establishing the usefulness of argumentation in structuring this kind of discussion, while the current case study instead serves to identify how disagreements about BNs are typically expressed and resolved manually by experts, and how this process can be specified and (partly) automated using formal (computational) argumentation. Such a formal specification allows us to investigate formal properties of our approach. Moreover, in future work our formal model can be the basis for developing and implementing software tools for supporting discussions about BNs between forensic experts, and for communicating their BNs and discussions to judges and prosecutors.

The paper is structured as follows. Section 2 provides some preliminaries on BNs and argumentation. In Section 3, we present our case study, on the basis of which we propose our argumentation-based approach to capturing and resolving conflicts about BNs in Section 4. In Section 5, we discuss related research and summarise our findings.

2 PRELIMINARIES

In this section, BNs and argumentation are reviewed.

2.1 Bayesian Networks

A BN [8] is a compact representation of a joint probability distribution $\Pr(\mathbf{V})$ over a finite set of discrete random variables $\mathbf{V}$. The variables are represented as nodes in a directed acyclic graph $G = (\mathbf{V}, \mathbf{E})$, where $\mathbf{E} \subseteq \mathbf{V} \times \mathbf{V}$ is a set of directed arcs $V_i \rightarrow V_j$ from parent V_i to child V_j. Each node V describes a number of mutually exclusive and exhaustive values; we will refer to these values as the *value space* of V. The BN further includes, for each node, a conditional probability table (CPT) specifying the conditional probabilities, or *probability parameters*, of the values of the node conditioned on the possible joint value combinations of its parents. A node is called *instantiated* iff it is set to a specific value. Given a set of instantiations, or evidence, the probability distributions over the other nodes in the network can be updated using probability calculus [8]. An example of a BN graph is depicted in Figure 1, where ovals denote nodes and instantiated nodes are shaded.

The BN graph G captures the independence relation among its variables. Let a *chain* be defined as a sequence of distinct nodes and arcs in the BN graph. A node V is called a *head-to-head node* on a chain c if it has two incoming arcs on c. A chain c is *blocked* iff it includes a node V such that (1) V is an uninstantiated head-to-head node on c without instantiated descendants; or (2) V is instantiated and has at most one incoming arc on c. A chain is inactive if it is blocked; otherwise it is called *active*. If no active chains exist

between V_1 and V_2 given instantiations of nodes in $\mathbf{Z}$, then they are considered conditionally independent given $\mathbf{Z}$.

2.2 Argumentation

Throughout this paper, a simplified version of the ASPIC$^+$ framework for structured argumentation [5] is assumed. We only semi-formally specify ASPIC$^+$ in this section, where the defined concepts will become clear through the examples discussed throughout this paper; for the formal definitions, we refer the reader to [5]. The ASPIC$^+$ instance assumes a *logical language* $\mathcal{L}$ containing the basic elements that can be argued about, a *knowledge base* $\mathcal{K} \subseteq \mathcal{L}$ of premises and a set of *inference rules* $\mathcal{R}$ that can be chained into arguments. Specifically, $\mathcal{L}$ is a non-empty propositional language with ordinary negation symbol $\neg$. Two propositions are each other's contradictories if one of them negates the other or if they are explicitly declared as contradictories by a contrariness function over the language (for example, being a bachelor and being married can be declared contradictories). Knowledge base $\mathcal{K}$ consists of two disjoint subsets $\mathcal{K}_n$ and $\mathcal{K}_p$ of axiom and ordinary premises, respectively. Inference rules in $\mathcal{R}$ are defined over $\mathcal{L}$ and are either strict or defeasible in that the consequent of the rule holds either deductively or tentatively given its antecedents. Arguments are then iteratively constructed from knowledge base $\mathcal{K}$ by chaining inference rules in $\mathcal{R}$.

Arguments can be *attacked* in three ways: on the conclusion of a defeasible rule (*rebuttal*), on a defeasible rule itself (*undercutting attack*), or on an ordinary premise (*undermining attack*). The conclusion of a strict rule or a strict rule itself cannot be attacked. To allow for undercutting attacks, $\mathcal{L}$ contains a well-formed formula corresponding to the name of every defeasible rule. Defeasible rules can be used to express argument schemes [9], which capture stereotypical non-deductive patterns of argumentation in a given domain as a scheme with a set of premises and a conclusion, plus a set of critical questions that need to be answered before the scheme can be used to derive conclusions. Critical questions then correspond to undercutters of the rule. Examples of argument schemes are presented in Section 3.2.

In ASPIC$^+$, attack is resolved into *defeat* on the basis of a binary *preference relation* over the arguments, where undercutting attack is preference-independent. ASPIC$^+$ thus induces abstract argumentation frameworks [10], which consist of a set of arguments along with a binary defeat[1] relation, and which can be used to evaluate the acceptability of the arguments. The theory of abstract argumentation frameworks is built around the notion of an *extension*, which is a set of arguments that is internally coherent and defends itself against attack. Dung precisely defined different types of extensions and conditions under which arguments belong to these extensions. For our current purposes, we only define the terms *"justified"*, *"overruled"*, and *"defensible"* argument, which are based on the notion of an extension. Informally, an argument is justified if it survived the competition, while an argument is called defensible if it is involved in a tie. An argument is justified iff it is a member of all extensions. An argument is overruled iff it is not justified and it is defeated by an argument that is justified. An argument

---

[1]Note that Dung [10] called his relation "attack". In ASPIC$^+$, the term attack is reserved for the basic notion of conflict, which is resolved into defeat using preferences.

is defensible iff it is not justified and not overruled. The status of a conclusion $\phi \in \mathcal{L}$ of an argument is then defined as follows. ϕ is justified iff there exists a justified argument with conclusion ϕ, ϕ is defensible iff ϕ is not justified and there exists a defensible argument with conclusion ϕ, and ϕ is overruled iff ϕ is not justified or defensible and all arguments with conclusion ϕ are overruled.

Throughout this paper, arguments are depicted informally in *argument graphs*; an example of an argument graph is depicted in Figure 2. Nodes in argument graphs denote propositions $\phi \in \mathcal{L}$, where shaded nodes denote propositions in $\mathcal{K}_n$. Every application of a rule in $\mathcal{R}$ is indicated with a solid (hyper)arc in the graph that is directed from the nodes corresponding to the antecedents of the rule to the node corresponding to the consequent of the rule. Every arc is annotated with the name of the applied inference rule (where d_i and s_i indicate a defeasible and a strict rule, respectively). A dashed arc directed from node N_1 to N_2 indicates that N_1 rebuts N_2. A dashed hyperarc directed from node $\neg d_1$ to a solid arc labelled d_1 indicates that $\neg d_1$ undercuts d_1.

3   CASE STUDY - DISAGREEMENTS ABOUT A FORENSIC FINGER MARK BN

In this section, we perform a case study to identify where disagreements about BNs can arise and how such disagreements are resolved by experts. Doekhie and colleagues [11, 12] constructed a BN for the interpretation of two finger marks[2], described in Section 3.1. Doshi [13] criticised this BN and proposed adjustments that address the supposed shortcomings. In Section 3.2, we analyse the argumentation structure of Doekhie's modelling decisions and Doshi's criticism on these decisions. While the main objective of this case study is to identify where disagreements about BNs typically occur and how such disagreements are resolved by experts, we also analyse the extent to which the constructed arguments can be classified as instances of existing or newly proposed argument schemes or as applications of critical questions of these schemes.

3.1   BN for Two Finger Marks at Finger Level

The BN that Doekhie and colleagues [11, 12] constructed is used to evaluate from which fingers two marks found at a crime scene originated, where the assumption is made that these marks are left behind by two consecutive fingers of the same hand in the act of a single touch. Specifically, if the fingers on the hands of a person are labelled 1 through 10, then the BN is used to assign a (posterior) probability that the two marks originated from a specific configuration of consecutive fingers. In a police investigation, a fingerprint expert enters the marks into a software system that automatically compares them to a fingerprint database. Knowing beforehand from which finger a mark most probably originated can considerably narrow down the search and, in turn, speed up the matching process. It should be noted that the constructed BN cannot be used to evaluate from which person the finger marks originated; BNs at person level are used for this purpose.

In Figure 1, Doekhie and colleagues' BN is depicted. FINGER-COMBINATIONS is the variable of interest for which we wish to obtain a posterior distribution. This variable describes eight values,


Figure 1: BN for the interpretation of two finger marks at finger level, constructed by Doekhie and colleagues [11, 12].

corresponding to the eight possible configurations of fingers from which the two marks originated. Specifically, these values are *1 and 2, . . . , 4 and 5, 6 and 7, . . . , 9 and 10*, where the first number denotes the finger number from which finger mark *A* originated and the second number indicates the finger number from which finger mark *B* originated. The FINGER A and FINGER B variables themselves describe ten values, corresponding to the ten possible fingers from which a mark can originate.

The GENERALPATTERN A and GENERALPATTERN B variables each describe a number of different values corresponding to the general patterns which are typically observed in finger marks and fingerprints, such as loops, whorls and arches. When using the BN in practice in a given case, these variables are instantiated to the general patterns observed in marks *A* and *B* to obtain a posterior distribution over the FINGERCOMBINATIONS variable.

The HAND variable accounts for the hand from which the two marks originated. This variable describes two values, namely *Left hand* and *Right hand*. The GENDER variable accounts for the gender of the donor from which the two marks originated. According to Doekhie and colleagues, this variable describes three values: *Male*, *Female* and *Unlabelled*. The CPT for this variable is filled using frequency statistics obtained from a *fingerprint database (D_1)*. This database contains information on every subject's gender and the finger from which each subject's print originated, as well as the general pattern of each fingerprint as labelled by a fingerprint examiner. In some cases, information regarding gender was not documented in the database and the prints were classified as "unlabelled". Frequency statistics from D_1 are also used to fill the CPTs for the GENERALPATTERN A and GENERALPATTERN B variables.

The CPTs for the HAND, FINGER A and FINGER B variables are filled using frequency statistics from a *finger mark database (D_2)*. This database contains similar information as D_1, except that the data is obtained from a large number of finger marks found at crime scenes instead of from fingerprints. As noted by Haraksim and colleagues [12], frequency statistics provide for a more informed prior than a uniform prior, which assigns equal prior probabilities to each finger or hand. From D_2, it can, for instance, be seen that marks of the thumb and index finger are found more often on crime scenes than marks originating from other fingers.

The CPT for the FINGERCOMBINATIONS variable is filled using frequency statistics from D_1. The exact manner in which these frequencies are chosen by Doekhie is not further discussed in this paper, as the parameterisation of the FINGERCOMBINATIONS variable is not criticised by Doshi.

---

[2]A finger mark is a mark recovered from a crime scene. A fingerprint is a print taken from a suspect at the police station.
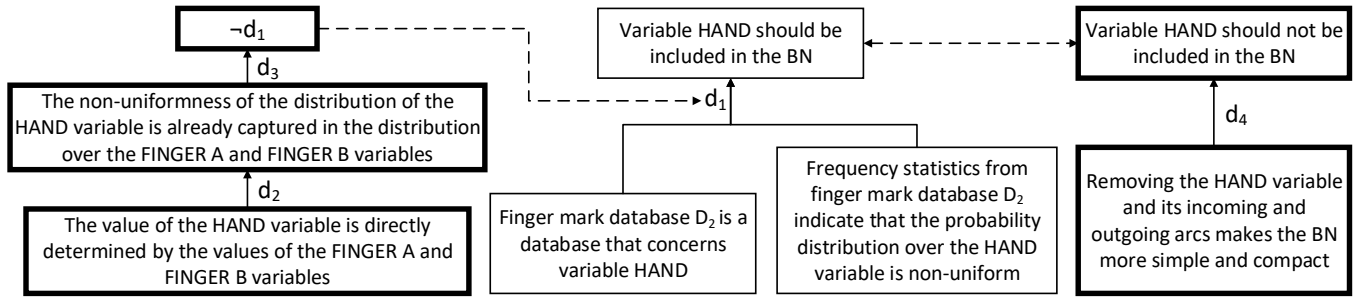
Figure 2: Argument-based analysis of Doekhie's modelling decision for including the Hand variable and Doshi's criticism.

3.2 Doshi's Criticism on Doekhie's BN

In this section, we analyse the argumentation structure of Doekhie's modelling decisions (Section 3.1) and Doshi's criticism on these decisions. In the argument graphs depicted throughout this section, propositions corresponding to Doekhie and colleagues' claims are indicated by plain boxes and propositions corresponding to Doshi's claims are indicated by thick boxes.

3.2.1 Relevance of the Hand Variable. Doekhie and colleagues [12, p. 102] state that frequency statistics from D_2 indicate that the probability distribution over the Hand variable is non-uniform, from which they conclude that this variable should be included in the BN and that this non-uniform distribution should be used for the parameterisation of this variable. This can be interpreted as an *argument from data set* (defined below): as D_2 implies a property of the Hand variable, a probability distribution (or BN) over a set of variables including the Hand variable should be constrained by this property. Doekhie and colleagues' arguments are depicted in the centre of Figure 2 (inference d_1).

The argument scheme from data set we propose is a generalisation of the scheme originally proposed by Keppens [6, p. 260]. Keppens' scheme can only be used to reason about the source of a specific parameterisation. We generalise this scheme such that it can also be used to reason about general properties of a BN, such as whether a variable should be included in the BN or not or what its value space should be:

S is a data set that includes variable(s) $V_1, \ldots, V_n$.
S implies a property *Prop* of a subset **V** of $V_1, \ldots, V_n$.

Therefore, a BN with variables **V** may be constrained by property *Prop*.

Critical questions corresponding to this scheme include (adapted from Keppens [6, p. 261]): (1) (Relevance) Does data set *S* cover all variables and values of variables necessary to identify the relevant circumstances covered by property *Prop*?; (2) (Representativeness) Is the population considered in data set *S* representative for the population under investigation in the present case?; (3) (Precision) Is the volume and precision of data set *S* consistent with the precision of property *Prop*?; and (4) (Consistency) Is the observation of property *Prop* in data set *S* consistent with other data sets?

Doshi criticises Doekhie and colleagues' modelling decision by stating that the value of the Hand variable is directly determined by the value of the Finger A and Finger B variables. Specifically, knowing from which finger a mark originated implies that we know from which hand the mark originated. From this claim it follows that the non-uniformness of the distribution of the Hand variable is already captured in the distribution over the Finger A and Finger B variables. This can be interpreted as an undercutter of the argument from data set as posed by Doekhie and colleagues; it is depicted on the left-hand side of Figure 2. We note that this undercutter cannot be categorised as any of the critical questions as proposed by Keppens. Instead it can be considered a variant of critical question (1), where the word "necessary" is replaced by the word "sufficient". We denote this critical question by (1′).

Doshi provides additional reasons for not including the Hand variable by claiming that this makes the BN more simple and compact. To capture this, we propose the following *argument scheme for reduced complexity*:

V is a variable in BN **B**.
Removing variable *V* and its incident arcs from **B** makes **B** less computationally and representationally complex.

Therefore, variable *V* should be removed from **B**.

We propose the following corresponding critical questions for this scheme: (1) Is variable *V* relevant for computing the posterior distribution of variables of interest in **B**?; and (2) Does the complexity gain from removing *V* compensate for the loss of accuracy and completeness of **B**?

Doshi's arguments are depicted on the right-hand side of Figure 2. Weighing the reasons pro and con, Doshi concludes that the Hand variable along with its incident arcs should be removed.

3.2.2 Independence of the GeneralPattern A and GeneralPattern B Variables. On discussing this with fingerprint examiners, Doshi found that these examiners believe that there exists no significant correlation/dependency between the general patterns that exist on different fingers. Doshi criticises Doekhie's BN by stating that the GeneralPattern A and GeneralPattern B variables are possibly dependent in her BN. The fingerprint examiners' claim that the GeneralPattern A and GeneralPattern B variables are independent can be interpreted as an *argument from expert opinion*. The argument scheme for arguments from expert opinion which we present below is taken from Prakken [7], who adapted it from Walton and colleagues [9]:
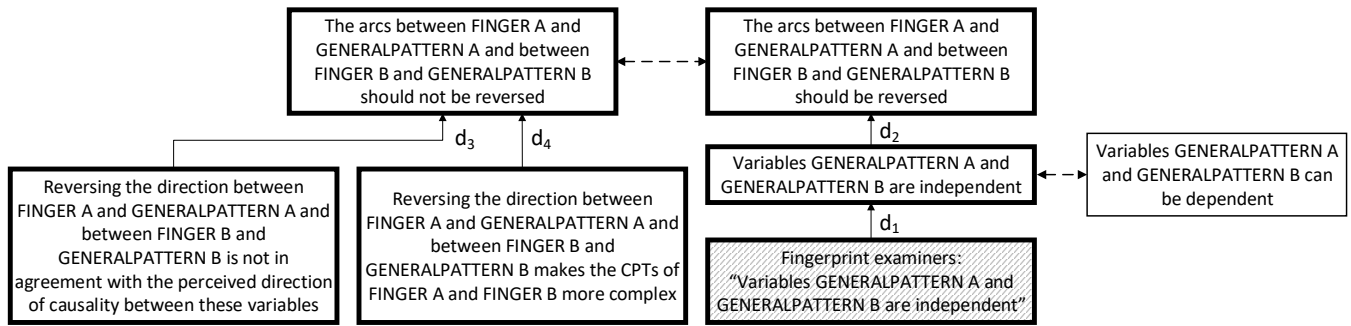
Figure 3: Argument-based analysis of Doshi's criticism on the possible dependency between the GENERALPATTERN A and GENERALPATTERN B variables.

E is an expert in domain D.
E asserts that P.
P is within D.

Therefore, presumably, *P*.

This scheme includes the following critical questions: (1) How credible is *E* as an expert source?; (2) Is *E* personally reliable as a source?; (3) Is *P* consistent with what other experts assert?; and (4) Is *E*'s assertion of *P* based on evidence?

The claim that the GENERALPATTERN A and GENERALPATTERN B variables are independent and the claim that these variables can be dependent rebut each other. Note that the latter claim is not explicitly made by Doekhie; it is, however, implied by the structure of the BN she constructed. Doekhie's implicit argument, as well as Doshi's counter-argument, are depicted in Figure 3.

Doshi notes that the GENERALPATTERN A and GENERALPATTERN B variables can be made independent in Doekhie's network by reversing the directions of arcs between GENERALPATTERN X and FINGER X, for both $X = A$ and $X = B$. However, this would (1) not be in agreement with the perceived direction of causality between these variables (inference d_3); and (2) would increase the number of probabilities that need to be estimated in the CPTs for the FINGER A and FINGER B variables and, therefore, make these CPTs more complex (inference d_4). The latter argument can be interpreted as an instance of a variation on the argument scheme for reduced complexity presented in Section 3.2.1. Instead of concerning the removal of a variable, this variation concerns the reversal of an arc between two variables. Specifically, if reversing an arc between two variables makes the BN less computationally and representationally complex, this arc should be reversed. We replace critical question (1) by "Does reversing arc *E* change the independence relation represented by the BN graph?". Critical question (2) can be directly applied to this scheme by replacing the words "removing *V*" by the words "reversing *E*".

Doshi's arguments for reversing the arc directions, as well as those for keeping the original arc directions, are depicted in Figure 3. These conclusions then rebut each other.

3.2.3 Conditional Independence of the FINGERCOMBINATIONS Variable and the FINGER A and FINGER B Variables. Doshi notes that, hypothetically, the FINGERCOMBINATIONS variable is conditionally

independent from the FINGER A and FINGER B variables given **Z** = {HAND, GENERALPATTERN A, GENERALPATTERN B} (its Markov blanket). Doshi criticises this modelling decision by stating that, knowing from which two fingers the marks originated, we should also be able to derive the combination of fingers that was used. Therefore, the FINGERCOMBINATIONS variable and the FINGER A and FINGER B variables should not be conditionally independent given **Z**. An argumentation-based analysis of Doshi's criticism is shown in Figure 4. Note that the argument on the left-hand side was not explicitly made by Doekhie and colleagues. Instead, it follows implicitly from the structure of Doekhie's BN. Its conclusion and the ultimate conclusion of Doshi's argument then rebut each other.

3.2.4 The Value Space of the GENDER Variable. Doekhie and colleagues base their design choice for using values *Male*, *Female* and *Unlabelled* as the value space of the GENDER variable on the observation that the GENDER variable can take on these three values in D_1. This can be interpreted as an argument from data set; it is depicted in the centre of Figure 5 (inference d_1). Doshi criticises this decision by arguing that in reality people are either male or female, and that "unlabelled" merely refers to the fact that data is missing with respect to this variable in D_1. Therefore (inference d_2), Doshi concludes that values *Male*, *Female* and *Unlabelled* are not mutually
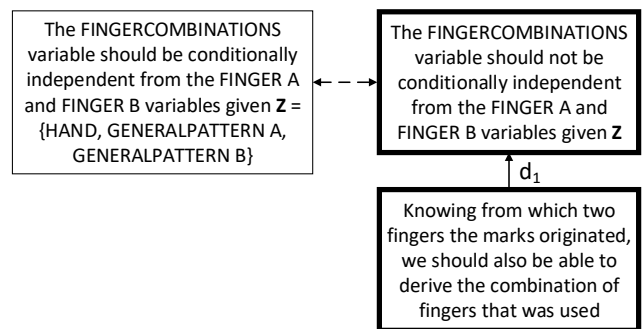
Figure 4: Argument-based analysis of Doshi's criticism on the independence between the FINGERCOMBINATIONS variable and the FINGER A and FINGER B variables given Z = {HAND, GENERALPATTERN A, GENERALPATTERN B}.
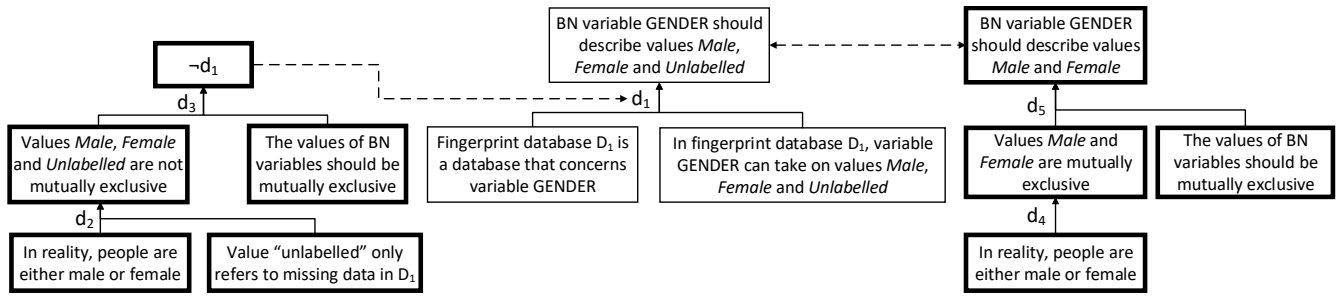
Figure 5: Argument-based analysis of Doekhie's modelling decision regarding GENDER's value space and Doshi's criticism.
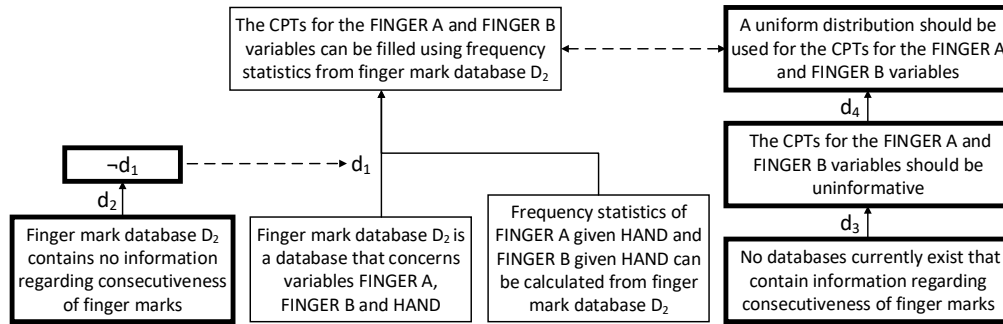

Figure 6: Argument-based analysis of Doekhie's modelling decision regarding the parameterisation of the FINGER A and FINGER B variables and Doshi's criticism.

exclusive. By furthermore stating that the values of variables in a BN should be mutually exclusive, Doshi poses (inference d_3) an exception to the argument from data set as posed by Doekhie and colleagues. Doshi's exception can be interpreted as an instance of critical question ($1'$) of this argument scheme; his arguments are depicted on the left-hand side of Figure 5.

In addition, Doshi claims that the GENDER variable should instead describe values *Male* and *Female*. First, he claims that values *Male* and *Female* are mutually exclusive, which is based on the fact that, in reality, people are either male or female. This argument is depicted on the right-hand side of Figure 5 (inference d_4). From this claim, and by again stating that the values of BN variables should be mutually exclusive, Doshi concludes that GENDER should describe values *Male* and *Female* (inference d_5). This argument and Doekhie and colleagues' argument then rebut each other.

Based on his criticism, Doshi proposed to adjust the BN of Doekhie and colleagues by removing the *Unlabelled* value from the value space of the GENDER variable.

3.2.5 Parameterisation of the FINGER A and FINGER B Variables. Doekhie's motivation for using frequency statistics from D_2 for the parameterisation of the FINGER A and FINGER B variables can be considered an argument from data set. This argument is depicted in the centre of Figure 6 (inference d_1). Doshi criticises this modelling decision by stating that D_2 cannot be used for establishing the relevant frequencies for these variables, as this database contains no information regarding consecutiveness of finger marks. This can be interpreted as an undercutter of Doekhie's argument from data set; specifically, it is an instance of critical question (1) of this

argument scheme. Doshi's argument is depicted on the left-hand side of Figure 6 (inference d_2).

Doshi instead proposes to use uniform distributions for both values of the HAND variable for the CPTs for these variables. His reason for using uniform distributions is that no databases currently exist that contain information regarding consecutiveness of finger marks and that the CPTs for the FINGER A and FINGER B variables should, therefore, be uninformative (for now). These arguments are depicted on the right-hand side of Figure 6; its ultimate conclusion and the conclusion of Doekhie's argument then rebut each other.

4 AN ARGUMENTATION-BASED APPROACH TO SUPPORTING DISCUSSIONS ABOUT BNS

In this section, we propose an argumentation-based approach that can be used to capture and help resolve conflicts about BNs. This approach is based on observations from our case study of Section 3. Our approach consists of two phases that are iteratively run through. In the first phase, domain experts are allowed to construct arguments pro and con the outcomes of modelling decisions underlying a given partially or fully specified BN, as well as preferences over these arguments. Arguments can be constructed regarding the inclusion or exclusion of different types of BN elements (described below) which may or not be in the existing BN. In the second phase, conflicts are then resolved by using the dialectical status of the constructed arguments to derive probabilistic and structural constraints on the BN (Sections 4.1 and 4.2).

Our sub-division of BN elements is inspired by our case study of Section 3, by the work of Pitchforth and Mengersen on the

validation of expert-elicited BNs [14] and by the work of Yet and colleagues [4]:

- Arguments regarding the existence of variables (Section 3.2.1). Conclusions of such arguments are of the form ϕ_V = "Include V in the BN" and $\neg\phi_V$ = "Exclude V from the BN".
- Arguments regarding whether a variable is observable or not. Typically, a fixed set of variables $\mathbf{O}$ is observed and instantiated when using the BN in practice in a given case. Conclusions of such arguments are of the form $\phi_{\mathbf{O}}^V$ = "Include V in $\mathbf{O}$" and $\neg\phi_{\mathbf{O}}^V$ = "Exclude V from $\mathbf{O}$".
- Arguments regarding (conditional) (in)dependencies (Sections 3.2.2 and 3.2.3). Arguments of this type are constructed for a given (possibly empty) subset of variables $\mathbf{Z} \subseteq \mathbf{V}$. Conclusions of such arguments are of the form $\phi_I^{V_1, V_2, \mathbf{Z}}$ = "Variables V_1 and V_2 should be conditionally independent given $\mathbf{Z}$" and $\neg\phi_I^{V_1, V_2, \mathbf{Z}}$ = "Variables V_1 and V_2 should be conditionally dependent given $\mathbf{Z}$".
- Arguments regarding the existence of arcs (Section 3.2.2). Conclusions of such arguments are of the form $\phi_E^{V_1 \rightarrow V_2}$ = "Include $E = V_1 \rightarrow V_2$ in the BN" and $\neg\phi_E^{V_1 \rightarrow V_2}$ = "Exclude $E = V_1 \rightarrow V_2$ from the BN". Arguments for arc reversal can then be indirectly constructed by constructing arguments with both conclusions $\neg\phi_E^{V_1 \rightarrow V_2}$ and $\phi_E^{V_2 \rightarrow V_1}$.
- Arguments regarding the value spaces of variables (Section 3.2.4). Conclusions of such arguments are of the form $\phi_{S_1}^V$ = "Use value space 1 for variable V". Such arguments may either be attacked by constructing an argument for its negation $\neg\phi_{S_1}^V$ = "Do not use value space 1 for variable V" or by constructing an argument for an alternative value space $\phi_{S_2}^V$ = "Use value space 2 for variable V", where $\phi_{S_1}^V$ and $\phi_{S_2}^V$ are declared contradictories of each other in $\mathcal{L}$.
- Arguments regarding the parameterisation of variables (Section 3.2.5). Conclusions of such arguments are denoted by ϕ_{Par}^V and are similar in form to arguments regarding value spaces of variables, and can be attacked similarly, by replacing the words "value space" by the word "parameterisation". Here, we assume that a conflict concerns *part* of the CPT for V. Specifically, let $V_1, \ldots, V_n$ be the parents of V. Then conflicts regarding the parameterisation of V concern a specific distribution $\Pr(V \mid V_1 = v_1, \ldots, V_n = v_n)$ for given values v_1 of $V_1, \ldots, v_n$ of V_n. We note that probabilities $\Pr(V = v \mid V_1 = v_1, \ldots, V_n = v_n)$ should sum to 1 when summing over all possible values v of V; conflicts regarding the parameterisation of a variable, therefore, never concern a single value of its CPT.

It should be noted that we do not propose a dialogue protocol. Instead, we first allow domain experts to construct arguments regarding all the different types of BN elements specified above along with a preference relation over those arguments, after which a priority ordering over the arguments corresponding to the different types of BN elements is applied to establish which conflicts, if any, should be resolved first.

4.1 Priority Ordering for Conflict Resolution

The idea of a priority ordering is based on the observation that BN elements from different types are dependent on one another.

For instance, it only makes sense to resolve conflicts regarding the value space of V iff V's existence is justified. We propose the following priority ordering over the types of BN elements:

P_1. Existence of variables.
P_2. Inclusion of variables in $\mathbf{O}$.
P_3. Conditional independencies between variables given $\mathbf{Z} = \mathbf{O}$.
P_4. Conditional independencies between variables given subsets $\mathbf{Z} \subseteq \mathbf{V}, \mathbf{Z} \neq \mathbf{O}$.
P_5. Existence of arcs between variables.
P_6. Value spaces of variables.
P_7. Parameterisation of variables.

A BN graph, by means of its arcs and their directions, represents an independence relation, so the (conditional) independencies implied by the graph should be verified. We note that, upon constructing BNs, arcs are typically added first, after which the implied independence relation is verified. In our approach, we instead prioritise the resolution of conflicts regarding independencies (P_3 and P_4) over those concerning arc directions (P_5), as arguments of the former type should generate new arguments for the existence of arcs and their directions at P_5 that together realise the conditional independence relation implied by the arguments at P_3 and P_4. Since an unjustified independence assumption can affect the behaviour of the BN, independencies exploited in deriving conclusions in an actual case are most important to verify. As such, we prioritise arguments regarding conditional independencies given $\mathbf{O}$ (P_3) over conditional independencies given subsets $\mathbf{Z} \subseteq \mathbf{V}, \mathbf{Z} \neq \mathbf{O}$ (P_4).

After obtaining a fully specified BN graph, conflicts regarding the value spaces of variables are considered (P_6). Finally, at P_7 conflicts regarding the parameterisation of variables are resolved. Specifically, the parameterisation of a variable V can be considered iff V's existence, V's value space, the existence of V's incoming arcs, and the value spaces of V's parents are justified.

We introduce the following notation. With $\mathcal{A}_{P_i}$ we denote the set of arguments whose conclusions concern BN elements at priority level P_i, $i \in \{1, \ldots, 7\}$. We assume that $\mathcal{A}_{P_i}$ divides into disjoint subsets of arguments $\mathcal{A}_{P_i}^e$ concerning BN element e at priority level P_i. For example, $\mathcal{A}_{P_1}^V$ only contains those arguments with conclusions ϕ_V and $\neg\phi_V$ for the inclusion/exclusion of variable V.

4.2 Constraint Table and Default Values

In Figure 7, the second phase of our approach, concerning conflict resolution, is summarised in pseudocode; throughout this section, we will explain and refer to the lines in this pseudocode. As input for this phase, either a fully or partially specified BN is used, along with a set of arguments and preferences constructed according to the first phase of our approach (lines 2 and 3 of the pseudocode). The process of conflict resolution at priority level P_i starts by calculating the dialectical status of conclusions in $\mathcal{A}_{P_i}^e$ for each BN element e (line 7 of the pseudocode). Conflicts are then resolved by consulting Table 1, which throughout this paper will be referred to as a *constraint table* (line 8 of the pseudocode). Depending on the priority level, different rows of the constraint table are used. At priority levels P_1 - P_5, the set of arguments $\mathcal{A}_{P_i}^e$ can either contain arguments with ultimate conclusions ϕ_e only, $\neg\phi_e$ only, or both. For these priority levels, rows 1−3 of the constraint table are used. At priority levels P_6

```
1  function ConflictResolution(BN, Args):
2      Input: Partially or fully specified BN
3      Input: Set of arguments Args with preferences
4      Output: Partially or fully specified BN
5      for i ∈ {1, . . . , 7} do
6          foreach BN element e at priority level P_i do
7              Calculate dialectical status of arguments in A^e_{P_i}
8              Consult constraint table:
9              if constraint table states "Include e" then
10                 Include e
11             end
12             else if constraint table states "Exclude e" then
13                 Exclude e
14                 Retain documentation regarding e
15                 Disregard arguments in A^e_{P_j} for j > i
16             end
17             else if constraint table states "No constraint" then
18                 Ask expert to further specify A^e_{P_i}:
19                 if A^e_{P_i} is further specified then
20                     goto line 7
21                 end
22                 else
23                     Resort to default choice
24                 end
25             end
26         end
27     end
28 end
```

Figure 7: Pseudocode of our argumentation-based approach to resolving conflicts about BNs.

and P_7, claims can also be constructed for alternative value spaces or parameterisations, indicated by $\phi_{e'}$; row 4 is also used at these priority levels.

The second column of the constraint table indicates the possible configurations of dialectical status of the claims in the first column. The third column then indicates the corresponding constraints on the BN for each configuration. Entries indicated with an asterisk indicate BN constraints at priority levels P_6 and P_7, whereas for the same dialectical status configurations at priority levels $P_1 - P_5$, entries without an asterisk are used.

Depending on the entry in the third column of the table, different actions should be taken. If the entry in the constraint table reads "Include e" or "Exclude e", the BN element should be included in the BN, respectively excluded (lines 9-11 respectively 12-16 of the pseudocode). In the latter case, the arguments regarding this variable should still be retained for documentation purposes in order to keep a "chain-of-evidence" of the changes made to the BN. If a domain expert wishes to provide further argumentation as to why this element should still be included, they can review the existing arguments regarding this BN element and supplement it with additional arguments or preferences. Furthermore, arguments of inferior priority which concern the excluded BN element should be disregarded in that these arguments should not be evaluated or

used in conflict resolution. For instance, if a variable V is removed from the BN at priority level P_1, arguments in $\mathcal{A}_{P_j}$ for $j > 1$ that concern this variable should be disregarded. Arguments for a conditional independency at P_3 and P_4 should be disregarded in case V is an element of the set of variables $\mathbf{Z} \subseteq \mathbf{V}$ under consideration or if V takes on the role of V_1 or V_2 in the argument.

For some configurations of dialectical status, a (univocal) constraint on the BN cannot be derived. Considering priority levels $P_1 - P_5$, two BNs are then possible: one in which e is included and one in which e is excluded. One possible approach would be to retain two BNs at this point, one including e and one excluding e, and to continue resolving conflicts for both BNs by considering all BN elements at all priority levels. Following this approach, a list (or tree) of candidate BNs is obtained after attempting to resolve all conflicts, from which the domain expert can choose one BN to continue with. However, for larger BNs the space of candidate BNs corresponding to a set of arguments would quickly become large. Moreover, if a large list of candidate BNs is provided to the domain expert, the differences between these BNs also need to be explained and presented to the domain expert in order for him to understand the differences between the presented BNs and to choose, which would require additional machinery.

We opt for an alternative approach in which the domain expert is asked to further specify his/her arguments every time a univocal constraint cannot be derived (line 18 of the pseudocode). Specifically, the domain expert is asked to provide additional arguments to supplement $\mathcal{A}_{P_i}^e$, or (further) specify a preference relation over $\mathcal{A}_{P_i}^e$. In case the expert opts to not further specify his arguments, a default choice is made by the approach (lines 22-24 of the pseudocode). What this default choice entails differs per priority level and per configuration of dialectical status. We note that, informally, the more complete the specification of the arguments and

Table 1: Constraint table for including BN element e. Entries indicated with an asterisk indicate BN constraints at priority levels P_6 and P_7, whereas for the same dialectical status configurations at priority levels $P_1 - P_5$ entries without an asterisk are used.

Proposed Claims	Dialectical Status	BN constraint
ϕ_e and $\neg\phi_e$	ϕ_e justified, $\neg\phi_e$ overruled	Include e
	ϕ_e overruled, $\neg\phi_e$ justified	Exclude e / No constraint*
	ϕ_e defensible, $\neg\phi_e$ defensible	No constraint
	ϕ_e overruled, $\neg\phi_e$ overruled	Further specification needed
ϕ_e only	ϕ_e justified	Include e
	ϕ_e overruled	Exclude e / No constraint*
	ϕ_e defensible	No constraint
$\neg\phi_e$ only	$\neg\phi_e$ justified	Exclude e / No constraint*
	$\neg\phi_e$ overruled	Include e
	$\neg\phi_e$ defensible	No constraint
ϕ_e and $\phi_{e'}$	ϕ_e justified, $\phi_{e'}$ overruled	Include e
	ϕ_e overruled, $\phi_{e'}$ justified	Include e'
	ϕ_e defensible, $\phi_{e'}$ defensible	No constraint
	ϕ_e overruled, $\phi_{e'}$ overruled	No constraint

the preference relation, the less the approach will resort to making a default choice. In the next subsections, the different configurations of dialectical status per priority level and their corresponding constraints are discussed in more detail.

In case the expert further specifies his arguments, he can recalculate the dialectical status of ϕ_e and/or $\neg\phi_e$ and establish the corresponding constraints. This process reiterates until BN element e is either included or excluded, or a default choice is made.

Conflicts About Variables. At priority level P_1, conflicts regarding the existence of variables are considered. In case claim ϕ_V for including variable V is justified (and $\neg\phi_V$ is overruled), this variable is included in the BN graph. If $\neg\phi_V$ is justified (and ϕ_V is overruled), variable V and all its incidents arcs are excluded from the BN graph. If ϕ_V and/or $\neg\phi_V$ is defensible, the default choice is to resort to using the original modelling decision, i.e. if variable V was included/excluded in the BN, then it should be included/excluded in the resulting BN. Lastly, if both ϕ_e and $\neg\phi_e$ are overruled, neither the choice for including nor excluding the BN element as a default is warranted. In this case, the expert should further specify his arguments to be able to derive a constraint (entry "Further specification needed" in the constraint table).

As an example of an application of our approach to resolving conflicts about variables, we consider the argumentation-based analysis presented in Section 3.2.1. In Figure 2, Doekhie's argument for including the HAND variable in the BN is undercut and is, therefore, overruled, as undercutting attack is preference-independent. It follows that Doshi's argument for excluding this variable is justified, as its only attacker is overruled. From the constraint table, it therefore follows that the HAND variable should not be included in the BN. In case we consider the set of arguments excluding Doshi's undercutter, both conclusions pro and con inclusion of the variable are defensible and no univocal constraint can be derived in this case. Doshi could in this case, for instance, specify that his argument for exclusion of the HAND variable is strictly preferred over the argument for inclusion of this variable. The dialectical status of the corresponding arguments is then recalculated; the argument for excluding the HAND variable is now justified, from which a constraint on the BN to exclude this variable is derived.

After resolving all conflicts at priority level P_1, conflicts at P_2, regarding the set of observable variables $\mathbf{O}$, are considered. Here, the terms "Include V" and "Exclude V"' now refer to including V in and excluding V from the set $\mathbf{O}$ rather than the BN graph.

Conflicts About Independencies and Arcs. At priority levels P_3 and P_4, conflicts regarding (conditional) independencies between variables are resolved using rows $1-3$ of the constraint table. In case $\phi_I^{V_1,V_2,\mathbf{Z}}$ is justified (and $\neg\phi_I^{V_1,V_2,\mathbf{Z}}$ is overruled), independence constraint I should be included in the set of independence constraints $\mathbf{I}$ to which the BN should adhere to. If $\neg\phi_I^{V_1,V_2,\mathbf{Z}}$ is justified (and $\phi_I^{V_1,V_2,\mathbf{Z}}$ is overruled), this independence constraint should not be included in $\mathbf{I}$. If $\phi_I^{V_1,V_2,\mathbf{Z}}$ and/or $\neg\phi_I^{V_1,V_2,\mathbf{Z}}$ is defensible, the default choice is to resort to the modelling decision implied by the original BN, that is, if V_1 and V_2 are conditionally independent given $\mathbf{Z}$ then I should be included in $\mathbf{I}$ and otherwise I can be excluded.

As an example, we consider the argumentation-based analysis presented in Section 3.2.3. The arguments depicted in Figure 4

concern priority level P_4, as Doshi considers the conditional independencies represented by the BN graph for a hypothetical set of variables $\mathbf{Z} = \{\text{HAND}, \text{GENERALPATTERN A}, \text{GENERALPATTERN B}\}$. As Doshi does not express an explicit preference, both rebutting conclusions in Figure 4 are defensible and no univocal constraint can be derived. In case Doshi expresses a strict preference for either of these conclusions, this chosen conclusion will be justified and the other overruled, in which case a constraint can be derived. We note that when using the BN of Doekhie and colleagues in an actual case, typically only the GENERALPATTERN A and GENERALPATTERN B variables are observed and instantiated, while the HAND variable is typically not instantiated. Doekhie and colleagues could, therefore, possibly counter Doshi's arguments at P_4 by constructing arguments for $\mathbf{O} = \{\text{GENERALPATTERN A}, \text{GENERALPATTERN B}\}$ at priority level P_2 and arguments at P_3 for a conditional dependency between the FINGER A and FINGER B variables and the FINGER-COMBINATIONS variable given $\mathbf{O}$. These arguments would then be prioritised over Doshi's arguments.

After resolving all conflicts at P_3 and P_4, our approach should generate new arguments for the existence of arcs and their directions at P_5 that together realise the conditional independence relation implied by $\mathbf{I}$. The precise manner in which this can be achieved is outside the scope of this paper; we note that so-called structure learning algorithms exist that can learn the structure of a BN graph given a set of independence constraints [8, chapter 7]. For now, we illustrate this part of the approach by means of an example. We consider the argumentation-based analysis presented in Section 3.2.2. In Figure 3, both arguments are constructed for the claims that GENERALPATTERN A and GENERALPATTERN B should be independent and can be dependent. Assuming Doshi strictly prefers his argument over Doekhie's implicit argument, his argument is justified. The corresponding independence constraint is, therefore, added to the set of independence constraints the BN should adhere to. Doshi then proceeds by attempting to find a configuration of arc directions under which GENERALPATTERN A and GENERALPATTERN B are independent, which generates arguments at P_5. The generated arguments at P_5 and the arguments that were already proposed by Doshi are then collectively considered, where conflicts at P_5 are resolved using rows $1-3$ of the constraint table.

Conflicts About Value Spaces. To resolve conflicts at priority level P_6, all rows of the constraint table are used. This table can only be used if arguments for at most two different value spaces for a variable V are constructed; similar constraint tables can be constructed in case more than two value spaces are proposed. In rows $1-3$ of Table 1, for entries in which ϕ_S^V is justified, value space S should be used for variable V. In case ϕ_S^V and/or $\neg\phi_S^V$ is defensible, no univocal constraint can be derived. In this case, the default is to resort to the value space of this variable as specified in the original BN. In case ϕ_S^V is overruled (and $\neg\phi_S^V$ is justified), no constraint can be derived, other than that value space S should not be used. In this case, the default choice is to resort to using a Boolean variable. However, if ϕ_S^V already concerns a Boolean value space, the expert should always further specify his arguments.

In the fourth row of Table 1, the situation is considered in which arguments are constructed for both $\phi_{S_1}^V$ and $\phi_{S_2}^V$. In case one of them is justified, the other is overruled; this results in the choice of

the value space posed in the conclusion of the justified argument. When $\phi_{S_1}^V$ and $\phi_{S_2}^V$ are both defensible, the default is to resort to the value space of this variable as specified in the original BN. If both $\phi_{S_1}^V$ and $\phi_{S_2}^V$ are overruled, neither value space should be used. If possible, the default is to again resort to a Boolean value space.

As an example of our approach to resolving conflicts about value spaces, we consider the argumentation-based analysis presented in Section 3.2.4. In Figure 5, arguments for two different value spaces are proposed, where the argument for Doekhie's value space is overruled as it is undercut. It follows that the argument for Doshi's value space is justified, as its only attacker is overruled. From the fourth row of the constraint table, it therefore follows that Doshi's value space should be used.

Conflicts About Parameterisations. Conflicts at priority level P_7 are resolved similarly as conflicts at P_6, where the default is to resort to using a uniform distribution in case arguments for specific parameterisations are overruled and the arguments are not further specified. Although a uniform distribution can easily be criticised, we opt for this distribution as a more informative choice cannot be made on the basis of the provided arguments. The default is also to resort to a uniform distribution even if the original distribution was uniform, instead of always requiring the expert to further specify his arguments in this case. Arguments in $\mathcal{A}_{P_7}^V$ regarding the parameterisation of a variable V should be disregarded in conflict resolution if one of V's incoming arcs, V's value space or one of V's parents' value spaces is overruled at P_5 or P_6. In this case, the default choice is to resort to a uniform distribution for V.

As an example, we consider the argumentation-based analysis presented in Section 3.2.5. In Figure 6, arguments for two different parameterisations are proposed. Similar to the example discussed in the previous paragraph, Doshi's parameterisation is justified and Doekhie's parameterisation is overruled; therefore, Doshi's parameterisation should be used.

5 CONCLUSION AND RELATED RESEARCH

In this paper, we have proposed an approach to capturing and resolving conflicts about BNs using computational argumentation. Our approach allows domain experts to document their reasons pro and con modelling decisions and their preferences in a structured manner using argumentation. The dialectical status of the constructed arguments is then used to derive probabilistic and structural constraints on the BN. Our approach always returns a partially or fully specified BN. Starting with a fully specified BN, a fully specified BN is returned by resorting to default choices in case a univocal constraint cannot be derived and the arguments and preferences are not further specified. Our approach can possibly be extended in future research by making explicit to the expert in which ways the arguments/preference relation can be further specified to obtain a univocal constraint. This would be an application of research on so-called resolution semantics; see e.g. [15].

As noted in the introduction, the work of Keppens [6] and Yet and colleagues [4] is related to this paper. In other related work, Bex and Renooij [16] and Wieten and colleagues [17] also derive constraints on a BN given arguments. The focus of their work is different from the present paper in that their aim is not to use arguments to argue about BNs but to map information expressed

as structured arguments [5, 18] to a BN. In other related work, approaches for explaining the reasoning patterns captured in BNs in terms of argumentation were proposed by Timmer and colleagues [19] and Keppens [20]. However, these approaches do not allow for any argumentative discussion about the construction of BNs.

In this paper, we have primarily focused on precisely specifying our approach using formal (computational) argumentation. Such a specification allows us to prove formal properties of our approach in future work. Moreover, in future work our formal model can be the basis for implementing software tools for supporting discussions about BNs between forensic experts, and for communicating their BNs and discussions to judges and prosecutors. Such tools may then also be used in evaluating our approach in practice.

REFERENCES
[1] N. Fenton and M. Neil. *Risk Assessment and Decision Analysis with Bayesian Networks.* CRC Press, 2012.
[2] F. Taroni, C.G. Aitken, P. Garbolino, and A. Biedermann. *Bayesian Networks for Probabilistic Inference and Decision Analysis in Forensic Science.* Wiley, 2014.
[3] N. Fenton, M. Neil, and D. Berger. Bayes and the law. *Annual Review of Statistics and Its Application,* 3: 51–77, 2016.
[4] B. Yet, Z.B. Perkins, N.R.M. Tai, and D.W.R. Marsh. Clinical evidence framework for Bayesian networks. *Knowledge and Information Systems,* 50(1): 117–143, 2017.
[5] S. Modgil and H. Prakken. A general account of argumentation with preferences. *Artificial Intelligence,* 195: 316–397, 2013.
[6] J. Keppens. On modelling non-probabilistic uncertainty in the likelihood ratio approach to evidential reasoning. *Artificial Intelligence and Law,* 22(3): 239–290, 2014.
[7] H. Prakken. A new use case for argumentation support tools: supporting discussions of Bayesian analyses of complex criminal cases. *Artificial Intelligence and Law,* https://doi.org/10.1007/s10506-018-9235-z.
[8] F.V. Jensen and T.D. Nielsen. *Bayesian Networks and Decision Graphs.* Springer, 2nd ed., 2007.
[9] D. Walton, C. Reed, and F. Macagno. *Argumentation Schemes.* Cambridge University Press, 2008.
[10] P.M. Dung. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial Intelligence,* 77(2): 321–357, 1995.
[11] G. Doekhie. A Bayesian network for assigning probabilities on which finger left a mark. Master's thesis. University of Amsterdam, The Netherlands, 2012.
[12] R. Haraksim, D. Meuwly, G. Doekhie, P. Vergeer, and M. Sjerps. Assignment of the evidential value of a fingermark general pattern using a Bayesian network. In A. Brömme and C. Busch, eds., *Proceedings of the 12th International Conference of the Biometrics Special Interest Group,* volume 212, pages 99–109. GI, 2012.
[13] V. Doshi. Validation of Bayesian networks with a case study on fingerprint general patterns. Master's thesis. Utrecht University, The Netherlands, 2013.
[14] J. Pitchforth and K. Mengersen. A proposed validation framework for expert elicited Bayesian networks. *Expert Systems with Applications,* 40(1): 162–167, 2013.
[15] S. Modgil and H. Prakken. Resolutions in structured argumentation. In B. Verheij, S. Woltran, and S. Szeider, eds., *Computational Models of Argument: Proceedings of COMMA 2012,* volume 245, pages 310–321. IOS Press, 2012.
[16] F. Bex and S. Renooij. From arguments to constraints on a Bayesian network. In P. Baroni, T.F. Gordon, T. Scheffler, and M. Stede, eds., *Computational Models of Argument: Proceedings of COMMA 2016,* volume 287, pages 95–106. IOS press, 2016.
[17] R. Wieten, F. Bex, H. Prakken, and S. Renooij. Exploiting causality in constructing Bayesian networks from legal arguments. In M. Palmirani, ed., *Legal Knowledge and Information Systems. JURIX 2018: The Thirty-first Annual Conference,* volume 313, pages 151–160. IOS Press, 2018.
[18] F. Bex, S. Modgil, H. Prakken, and C.A. Reed. On logical specifications of the argument interchange format. *Journal of Logic and Computation.* 23(5): 951–989, 2013.
[19] S.T. Timmer, J.-J.C. Meyer, H. Prakken, S. Renooij, and B. Verheij. A two-phase method for extracting explanatory arguments from Bayesian networks. *International Journal of Approximate Reasoning,* 80: 475–494, 2017.
[20] J. Keppens. Argument diagram extraction from evidential Bayesian networks. *Artificial Intelligence and Law,* 20(2): 109–143, 2012.