



Utrecht University

MASTER'S THESIS

**Formulating guidelines for the
manual construction of BN
graphs in data-poor domains**

Author:
Tigris de Leeuw
4213866

Daily supervisor:
G. M. Wieten, MSc.

First examiner:
Dr. S. Renooij

Second examiner:
Prof. Dr. F. J. Bex

A thesis presented for the degree of MSc Artificial Intelligence

August, 2020

Abstract

Bayesian networks (BNs) are powerful mathematical tools that have found applications in many fields where uncertainty plays a role, such as medicine, forensics and law. In short, a BNs consist of a graph and a set of probability parameters, and can be used to compute a probability of interest. In certain domains, data is becoming increasingly available; hence, in these domains BNs can often be constructed directly from data. In data-poor domains however, the network has to be handcrafted with the help of a domain expert. The current literature provides little guidance on manual construction of the BN graph. In this thesis we present a set of guidelines and comparison measures that provide practical aid in building a BN graph by hand with the help of a domain expert for a problem from any domain. In contrast to earlier approaches, we aim for a set of guidelines that can be used for cases from any domain while still maintaining a high degree of practical applicability. To this end, we carry out a comprehensive literature study on current forms of guidance in this area, a detailed examination of advantageous characteristics for BN graphs, and create and assess a database of manually constructed BN graphs from the literature.

To evaluate the applicability of these guidelines and measures we carry out a case study, where we take BNs constructed in the forensic and legal domains as our case. We find that we largely succeeded in providing guidance for an inherently subjective process (i.e. manual BN graph construction). In addition, the BN graphs from the forensic and legal domains are generally in line with what our guidelines describe. Testing the applicability of our comparison measures however proved difficult, due to the fact that the literature provided few different BN graphs for the same case.

Contents

1	Introduction	4
1.1	Research questions	5
1.2	Outline	7
2	Preliminaries	8
2.1	Graph theory and probability theory	8
2.1.1	Graph theory	8
2.1.2	Probability theory	9
2.1.3	Conditional independence	9
2.2	Bayesian networks	10
2.2.1	Causality and intercausal interactions in a BN	11
2.3	Overview of the Bayesian network construction process	12
2.3.1	Obtaining the network parameters	13
2.3.2	Model validity	13
I	Guidance from the literature on manual BN graph construction	15
3	Manual graph construction – guidance from the literature	16
3.1	Identification of the variables and their values	16
3.2	Building the graph	17
3.2.1	The notion of causality	17
3.2.2	Other types of relationships	18
3.2.3	An idiom-based approach	19
3.3	A note on model complexity	23
4	BN construction in the forensic domain: A template by Taylor and colleagues	25
4.1	The evaluation of forensic traces	25
4.2	The use of BNs in the forensic domain	26
4.3	A template for activity level BN construction for biological traces	27
5	Advantageous properties for BN graphs constructed from arguments	29
II	Formulating new guidelines and comparison measures	31
6	Manual BN graph construction - Obtaining observations from the field	32

7	A new set of advantageous properties for BNs from any domain	33
7.1	Conclusion	40
8	Assessment of the template by Taylor and colleagues by means of the new properties	41
9	Guidelines	45
9.1	Identification and classification of the variables	47
9.2	Designing the graph	50
9.3	Evaluation of the structure	52
10	Comparison measures	54
10.1	Development of the measures	54
10.2	Final measures	55
10.3	Measure 1 - Computational complexity	56
10.4	Measure 2 - Representation	56
10.5	Priority ordering	56
III	Case Study	58
11	Evaluating the guidelines and comparison measures for the forensic and legal domains	58
11.1	Variable identification	59
11.2	Network structure	60
11.3	Evaluation of the structure	62
11.4	Comparison measures	63
12	Discussion	64
13	Conclusion	66
	Appendices	74
A	Figures of BNs	74
B	List of articles for the database of forensic and legal BNs	88

1 Introduction

The ubiquity of uncertainty and randomness complicates the decision-making involved in solving real-world problems. The handling of uncertainty therefore has a pivotal role in the field of artificial intelligence (AI): an intelligent system must be able to perform a task that requires intellect, based on its knowledge and environment, even when its knowledge and/or understanding of its environment is incomplete or uncertain. The predominant way of dealing with uncertainty is with probability theory. One way of interpreting probabilities is as relative frequencies of specific outcomes, but probabilities can also be used in a more subjective way: to summarise the extent to which something is believed. In this way, one can express a degree of uncertainty with probability theory [24]. Several tools exist that support domain experts in modelling (domain) knowledge whilst incorporating the uncertainties that go along with that, in order to perform probabilistic inference to reason about the probability of particular outcomes.

One such tool is a Bayesian Network (BN) [38]. BNs are powerful mathematical tools that have found applications in many different fields where uncertainty plays a role, such as medicine, forensics and law [24]. In short, a BN describes a joint probability distribution over a set of variables by means of a graph and a set of probability parameters.

For a given problem, there is no ‘one true BN model’: depending on the intentions and preferences of those involved in building the model, different models may be constructed varying in the structure of the graph and the specification of the network parameters (parameterisation). The construction of a BN is a complex, iterative process [13] in which the structure of the graph and corresponding parameters can be learnt from data, established manually, or found through a combination of these. A manual approach is used for cases from data-poor domains, such as forensics or law. For these domains, the graph and probabilities can be established by hand with help of an expert in the domain of the problem.

The graph of a BN captures the independence relation among its variables from the modelled problem, and specifies which probabilities have to be estimated. For BNs from data-poor domains, it is therefore essential that the graph correctly represents the beliefs of the domain expert in their view of the relevant variables and the correlations between them. The graph also functions as a visual communication tool, showcasing the variables and reasoning considered in the case.

Even though data-poor domains require a manual approach to BN graph construction, few methods or guidelines for handcrafting a BN graph exist. The aim of this thesis therefore is to provide aid for a BN engineer tasked with manually constructing the graph of the BN for a data-poor domain (i.e. with help of a domain expert). Specifically, we propose a set of generally applicable guidelines for the manual construction of BN graphs in data-poor domains, as well as a set of specific measures to compare BNs on. In doing so this study contributes to untangling the complex process of BN construction, such that

their application to problems from different data-poor domains can increase.

1.1 Research questions

The ultimate goal in this thesis is to provide a set of guidelines and comparison measures. This aim can be divided into two objectives. The first is to research what manually constructed BN graphs should ideally look like, such that we know what our guidelines aim to achieve. Secondly, we use this insight to formulate a set of guidelines for BN graph construction in data-poor domains. In addition, we propose a set of measures with which different BN structures for the same case can be assessed on. The newly proposed guidelines and measures should be evaluated on whether they are in line with practices from the field. The forensic and legal domains are data-poor domains where a significant amount of work is done on manual BN construction. These domains are treated as a case study, and the guidelines and measures are related to this ‘case’.

Concretely, this leads to the following research questions:

1. What are desirable characteristics for manually constructed BN graphs for problems from data-poor domains and why are these desirable?
2. How can these insights be combined to formulate:
 - (a) a set of practical guidelines for the manual BN graph construction with help of domain experts for a problem from any domain, and
 - (b) a set of comparison measures for analysing different BN graphs for the same problem and establishing which graph is the most suitable in a specific situation?
3. How do the guidelines and comparison measures from research question 2 relate to actual manually constructed BN graphs from the forensic and legal domains?

To answer the first research question we first provide a comprehensive overview of the guidelines on manual BN graph construction that are currently provided by the literature. Additionally, we gather information on the manual BN graph construction process by examining the characteristics of handcrafted structures from the field, specifically from the forensic and legal domains. The guidelines from the literature are examined on their applicability to domains in general, and on the characteristics of their resulting graphs. The goal is to gather inspiration for what the new guidelines and measures should look like, as well as what they should achieve. This insight is translated into the formulation of practical set of new guidelines in order to answer research question 2a. Where the methods from the literature might be domain-specific, our set of guidelines is general-purpose¹. We also use our research to draw up a set of comparison measures for the informed evaluation of BN graphs built for the same case, as specified in research question 2b. Although these measures are not as detailed as we had

¹General-purpose in that they are not domain-specific.

hoped, they provide a good first step towards a structured way of comparing the structure of BNs on particular aspects, which has not been done before. In the final part of this thesis we evaluate the applicability of the newly drafted guidelines and measures by carrying out a case study, where the ‘case’ is both the forensic and legal domains. We show that the guidelines are applicable, but argue that they can be tailor-made to accommodate to the specifics of particular domains.

The approach to realising this thesis’ objective can be divided into three parts:

- An extensive literature study to summarise current guidelines and methods for manual BN graph construction;
- A critical assessment of these guidelines and methods to gather inspiration for our own guidelines and comparison measures;
- Observations from the “field” to determine the practices and choices in BN graph construction.

In the first place we provide a thorough overview of the literature on manual BN graph construction. Different approaches to designing the structure of a BN, from general to domain-specific, are summarised. First of all, this serves as an overview of previous work on the subject. Secondly, a synopsis of current guidance on this topic is useful for several reasons. It gives a detailed overview of existing methods, providing a comprehensive summary of a the possible courses to take when designing a BN graph by hand. A literature overview such as we present in this thesis has not been done before: the methods we discuss illustrate both the scarcity of practical and general guidance for manual BN graph construction as well as the diversity in approaches to this topic. As such it contributes towards providing clarity and continuity to this part of the field.

Additionally, such an overview allows for a detailed assessment of the current methods for handcrafting the BN graph. In doing so we identified strengths and weaknesses in two of the methods from the literature, that provided us with the knowledge on how our own guidelines and measures should be drafted and what pitfalls to avoid. In addition to gathering inspiration, we present an evaluation of the strengths and shortcomings of these methods and a list of advantageous BN graph properties, two other contributions of this thesis.

Lastly, we gathered knowledge on manual BN graph construction on a more concrete level from observations from ‘the field’. We did so in two ways. First of all, a semi-structured interview with a person that is both a BN engineer and domain expert was conducted to gain insight into the BN graph construction process. Secondly, we studied manually constructed BNs from the forensic and legal literature. These domains were chosen because they are data-poor, similar in their approach to BN construction and provide a rich variety of handcrafted BNs. A database of manually constructed BN graphs from the literature was created and assessed on various aspects. To our knowledge, this is the first time a repository of this magnitude, solely consisting of BNs from the literature, has

been devised, which characterises a novel approach to researching handcrafted BN graph traits.

1.2 Outline

This thesis is structured as follows: first, Section 2 provides an overview of the basic concepts necessary to understand BNs, after which BNs are introduced. Part I consists of Sections 3, 4 and 5, which together form the literature study. Different forms of guidance are discussed, ranging from general construction approaches in Section 3 to a highly domain-specific but practical method in Section 4. In contrast to the stepwise method for BN graph construction from Section 4, Section 5 provides a list of properties that BN graphs are desired to possess. Parts II and III include the research we carried out for this thesis. In Section 4 we describe the method with which we have obtained knowledge of manual BN graph construction in practice. Desirable properties to which BN graphs should adhere are established in Section 7. As our ambition is to formulate practical guidelines, in Section 8 we examine the practicality of the domain-specific method from the literature (Section 4) by relating this method to the new properties as presented in Section 7. Following these insights, we propose our guidelines and measures in Sections 9 and 10. These are tested in a case study in Part III: here we relate the guidelines and measures to BNs from the forensic and legal domain. We discuss our research in Section 12 and conclude our findings in Section 13.

2 Preliminaries

In this section we provide preliminaries for the subsequent parts of this thesis. We start by introducing some concepts from graph theory in Section 2.1 which are relevant for Bayesian networks, the basics of which are introduced in Section 2.2 (see for example [24] for a more elaborate introduction). Since this thesis addresses the question of drawing up a set of guidelines for manual BN construction, the general process of building a BN is explained in Section 2.3. More specific guidance from the literature as to the development of the graphical aspect of a BN is not discussed in this section, but in subsequent Sections 3, 4 and 5.

2.1 Graph theory and probability theory

The relevant notions from graph- and probability theory are discussed in this section.

2.1.1 Graph theory

An *undirected graph* G is a pair $(\mathbf{V}, \mathbf{Ed})$, where \mathbf{V} is a finite set of nodes and \mathbf{Ed} is set of unordered pairs (V_i, V_j) , $V_i, V_j \in \mathbf{V}$, called *edges*. A *directed graph* (or digraph) G is a pair (\mathbf{V}, \mathbf{A}) , where \mathbf{V} is a finite set of nodes and \mathbf{A} is set of ordered pairs (V_i, V_j) , $V_i, V_j \in \mathbf{V}$, called *arcs* [21, p. 10]. An arc is denoted as $V_i \rightarrow V_j$, where V_i is the *tail* (or tail terminal) and V_j is the *head* (or head terminal) [19]. The *underlying graph* of a digraph is obtained by replacing each arc $(V_i, V_j) \in \mathbf{A}$ by an edge (V_i, V_j) , resulting in an undirected graph [21, p. 11].

A node $V_i \in \mathbf{V}$ in a digraph is considered a *parent* of another node V_j if $V_i \rightarrow V_j \in \mathbf{A}$. V_j in this case is considered the *child* of V_i . $\rho(V_i)$ is the set of all parents of node V_i , $\sigma(V_i)$ denotes its set of children. The reflexive, transitive closure of V_i under the parent relation, denoted as $\rho^*(V_i)$, is the set of ancestors of V_i . Similarly, $\sigma^*(V_i)$ denotes the set of descendants of V_i . $\sigma(V_i) \cup \rho(V_i)$ defines the set of *neighbours* of V_i , the size of which is its *degree*. The number of parents of V_i , i.e. the size of $\rho(V_i)$, is its *in-degree* [21, p. 10].

A *path* from V_0 to V_k in a digraph is a sequence of distinct nodes V_0, \dots, V_k , $k \geq 1$, with distinct arcs $V_{i-1} \rightarrow V_i \in \mathbf{A}$, $1 \leq i \leq k$, between them. k is the length of the path. A *chain* in a digraph is similar to a path, but does not take the arc direction into account. A *cycle* in a digraph is a path from a node V_i to V_i of non-zero length. In case a digraph contains a cycle it is called *cyclic*; otherwise, it is called *acyclic* [21, p. 11].

2.1.2 Probability theory

Let \mathbf{V}^2 be a set of n variables $\{V_1, \dots, V_n\}$. A variable $V_i \in \mathbf{V}$ can assume one out of a set of *values* (or states). A variable V can for example assume the values *true* and *false*, denoted as v and $\neg v$ respectively, making it binary-valued³.

A *joint probability distribution* \Pr over \mathbf{V} , written as $\Pr(\mathbf{V})$, describes the probability of the variables in \mathbf{V} taking on specific values. This probability ranges from 0 to 1, with 0 meaning ‘certain not to occur’ and 1 expressing ‘certain to occur’. The probabilities from the same (conditional) distribution for each state of a variable sum to 1, e.g. $\Pr(v_i) + \Pr(\neg v_i) = 1$. If a distribution requires m parameters to be estimated, $m - 1$ of these are *free* in that their probability can be chosen freely as long as their sum is less than or equal to one. The last remaining parameter is *bound*, because its probability must be the number needed to complete the sum to one [21, p. 9].

The *prior probability* describes the probability for a value in the absence of any other information. To illustrate, $\Pr(v_i) = 0.2$ denotes that there is a 20% chance that V_i will take on the value v_i if nothing else is known. If more information is available, i.e. if evidence is obtained, *conditional probabilities* are used to specify the (revised) probability of a variable taking a particular value. For two binary variables V_i and V_j , consider $\Pr(v_i | v_j)$: this expresses the probability of v_i given that v_j is known with certainty (and nothing else is known that might affect this probability).

2.1.3 Conditional independence

(Conditional) independence can be defined in terms of numerical quantities, i.e. the independence relation is captured by the probabilities for the distribution at hand. For disjoint sets of variables $\mathbf{V}_i, \mathbf{V}_j \in \mathbf{V}$, if $\Pr(\mathbf{V}_i | \mathbf{V}_j) = \Pr(\mathbf{V}_i)$, \mathbf{V}_i and \mathbf{V}_j are *independent*. Knowledge about \mathbf{V}_j does not influence the probabilities of \mathbf{V}_i . \mathbf{V}_i and \mathbf{V}_j are *conditionally independent* given $\mathbf{V}_k \in \mathbf{V}$ if $\Pr(\mathbf{V}_i | \mathbf{V}_j, \mathbf{V}_k) = \Pr(\mathbf{V}_i | \mathbf{V}_k)$. In other words, if \mathbf{V}_k is known, information of \mathbf{V}_j becomes irrelevant for \mathbf{V}_i .

An independence relation amongst variables can also be represented in graphs. We consider digraphs only, where the *d-separation criterion* can be used to derive conditional independencies from a digraph. First, we introduce the concept of *blocking*. Suppose three nodes $V_1, V_2, V_3 \in \mathbf{V}$ appear consecutively on a chain s . s is blocked by a (possibly empty) set of variables \mathbf{Z} if one of the following holds:

- arcs $V_1 \rightarrow V_2$ and $V_2 \rightarrow V_3$ form a serial connection on chain s , and $V_2 \in \mathbf{Z}$;

²Throughout this thesis we will use capital letters to denote a single variable V , which when boldfaced, i.e. \mathbf{V} , indicate a set of variables. Lower case is used for a specific value v of V . In case of a binary valued variable V with values *true* and *false*, we use v and $\neg v$ respectively.

³For simplicity, in this thesis we assume binary-valued variables to take on values *true* and *false*, unless stated otherwise.

- arcs $V_2 \rightarrow V_1$ and $V_2 \rightarrow V_3$ form a diverging connection on chain s , and $V_2 \in \mathbf{Z}$;
- arcs $V_1 \rightarrow V_2$ and $V_3 \rightarrow V_2$ form a converging connection on chain s , and $\sigma^*(V_2) \cap \mathbf{Z} = \emptyset$;

To illustrate, see Figures 1, 2, and 3, where the shaded nodes are part of the blocking set \mathbf{Z} .




Figure 1: A serial connection




Figure 2: A diverging connection




Figure 3: A converging connection

We now introduce the concept of *d-separation*: two nodes V_1 and V_2 are *d-separated* by a (possibly empty) set of nodes \mathbf{Z} if every chain between V_1 and V_2 is blocked by \mathbf{Z} . Two nodes V_1 and V_2 that are *d-separated* given \mathbf{Z} are considered *conditionally independent* given \mathbf{Z} [21, p.27].

Lastly we discuss the concept of *Markov equivalence* in DAGs. The *skeleton* of a DAG is the undirected underlying graph of the DAG. An *immorality* is a converging connection (see Figure 3) with two parents that are not directly connected by an arc. Two DAGs are *Markov equivalent* if they share the same skeleton and have the same immoralities [19].

2.2 Bayesian networks

A Bayesian network [24, 38] describes a joint probability distribution \Pr over a set of random variables \mathbf{V} . A BN consists of a qualitative and a quantitative part.

The qualitative part includes an acyclic digraph (DAG) G . This graph $G = (\mathbf{V}, \mathbf{A})$ represents the set of random variables \mathbf{V} as nodes and \mathbf{A} is a set of directed arcs $V_i \rightarrow V_j$, $V_i, V_j \in \mathbf{V}$, which describe influences between the variables⁴. The combination of nodes and arcs is often referred to as the *structure* or *graphical component* of a BN, and represents the independence relation of \mathbf{V} . The *d-separation criterion*, introduced in Section 2.1.3, can be used to derive conditional independencies from G .

Each node in a BN is associated with a set of (conditional) probabilities, summarised in a conditional probability table (CPT). These together make up the quantitative part of a BN. For a variable V , the CPT defines the conditional probabilities (network parameters) of the values of each node given each possible configuration of outcomes of its parents.

⁴Technically a node represents a variable, but throughout this thesis the terms ‘variable’ and ‘node’ are used interchangeably.

Figure 4 shows an example of a BN for computing the probability of a particular cause of a fire incident, constructed by Biedermann and colleagues [3]. All nodes are binary-valued, except node ‘H’ which has 3 states $\{h_1, h_2, h_3\}$. Only the ‘free’ parameters (see Section 2.1.2) are given. For a description of the nodes we refer the reader to Figure 23 in Appendix A.

Figure 4: The BN from Biedermann and colleagues [3]




Figure 5: The graph

$\Pr(h_1) = 0.0203$	$\Pr(m) = 0.95$
$\Pr(h_2) = 0.3452$	
	$\Pr(l \mid d, b) = 1$
$\Pr(d \mid m, h_1) = 0$	$\Pr(l \mid d, \neg b) = 1$
$\Pr(d \mid m, h_2) = 0$	$\Pr(l \mid \neg d, b) = 1$
$\Pr(d \mid m, h_3) = 0.02$	$\Pr(l \mid d, \neg b) = 0$
$\Pr(d \mid \neg m, h_1) = 0$	
$\Pr(d \mid \neg m, h_2) = 0$	$\Pr(x) = 0.75$
$\Pr(d \mid \neg m, h_3) = 0$	
	$\Pr(t \mid x, l) = 1$
$\Pr(b) = 0.05$	$\Pr(t \mid x, \neg l) = 1$
	$\Pr(t \mid \neg x, l) = 1$
$\Pr(e \mid t) = 0.99$	$\Pr(t \mid \neg x, \neg l) = 0$
$\Pr(e \mid \neg t) = 0.001$	

Table 1: The network parameters

With the network parameters, a joint probability distribution can be defined on \mathbf{V} which respects the independence relation defined by the structure, with the chain rule for Bayesian networks [24]:

$$\Pr(\mathbf{V}) = \prod_{i=1}^n \Pr(V_i \mid \rho(V_i)) \quad (1)$$

With this joint probability distribution, any (prior or posterior) probability of interest can be calculated from the network. Doing so is a complex process, that when approached without exploiting the independences represented by the structure quickly becomes computationally infeasible [24]. For this reason several algorithms have been proposed in the literature, such as that of Lauritzen and Spiegelhalter [33].

2.2.1 Causality and intercausal interactions in a BN

Although the only requirement for the qualitative component of a BN is that it encodes the independence relation of the variables from the modelled case,

the arcs in a BN graph often have a *causal* connotation. That is to say, an arc $V_i \rightarrow V_j$ signifies a relation where V_i is the cause of V_j . Directing arcs according to this notion of causality facilitates the process of correctly modelling the (in-)dependencies in the case (although these are not guaranteed by directing arcs in a causal manner) [26], and eases probability estimation (discussed further in Section 3.2.1).

The BN graph can serve as a way to represent *intercausal interactions*⁵. One such interaction is the *explaining-away* effect, which can occur when two variables V_i and V_j share a common child V_k (a converging connection, see Figure 3). Both V_i and V_j can explain the occurrence of V_k . After observing V_k or one of its descendants, an increase in the probability of V_i (V_j) can result in a lower probability for V_j (V_i), and hence ‘explains away’ V_j (V_i). In case a higher probability of one parent leads to a higher probability of the other, an effect takes place called *explaining-in* [54]. Other types of intercausal reasoning include noisy functional dependence (such as (leaky) noisy-or and noisy-and, see [24]) and intercausal cancellation (see for example [58]). The type of intercausal interaction is partially governed by the graph structure, but its occurrence is determined by the parameterisation.

2.3 Overview of the Bayesian network construction process

BNs can be constructed in different ways: the structure and parameters can be extracted automatically from data, constructed by hand with help of a domain expert, or established through a combination of these. Regardless of the approach, the construction of a BN generally involves two tasks: establishing the qualitative and the quantitative component of the BN model. The first task can be further divided into two sub-tasks, namely the identification of variables and their values, and the graph design. In short, BN construction involves:

1. Defining the BN graph:
 - (a) Identification of the relevant variables from the case to be modelled and specification of their values;
 - (b) Determining the graphical structure.
2. Specifying the network parameters (*parameterisation*).

As the structure determines which probabilities have to be estimated, the steps are initially carried out sequentially. Finalising a BN model however is an iterative process where engineers will go back to and alternate between the steps, assessing the resulting BNs in between, before ending up with a conclusive, satisfactory BN [13].

⁵We note that, while the term ‘*intercausal* interaction’ is used, these interactions can occur regardless of the type of relation between parents and child [12].

Regarding automated BN construction, there are several algorithms to establish the qualitative component (step 1 in the construction process). Constraint-based algorithms aim to identify a structure from data, where the conditional independence relationships amongst the variables serve as constraints for the structure. Score-based algorithms take a different approach: they determine how well a given structure fits the data for which it is built, according to some scoring criterion [8, 9]. For a given structure, parameters can be learnt or estimated from data (through for example Maximum Likelihood Estimation or Maximum A-Posteriori estimation).

As this thesis will deal with data-poor domains, we will not further consider automatic construction of BN models from data (but see for example [8] for a discussion on several algorithms for learning BN structures). Handcrafted BNs are built in close collaboration with a domain expert, who is involved in all steps described above. Sections 3.1 and 3.2 discuss in depth the steps involved in designing the qualitative component of the BN (i.e. steps 1a and 1b). As these are considered guidance from the literature on building the BN graph, we elaborate on them in Part I of this thesis.

2.3.1 Obtaining the network parameters

Step 2, i.e. establishing the quantitative component of the BN, involves estimating the network parameters for each node in the network. Often these estimates are taken from literature on the domain to which the BN belongs, from statistical data or elicited from domain experts [13]. Various techniques exist to obtain probabilities through expert elicitation, such as numerical probability scales or reference lotteries [40].

Nevertheless, these sources rarely present all necessary probabilities for the newly constructed model: there might be missing information, data can be biased, data might not exactly include the variables established in the network, and so on [21, p. 85]. Obtaining reliable estimations therefore is a difficult and time-consuming task [13].

Inaccurate probabilities negatively influence the reliability of the output of the BN model. Performing a *sensitivity analysis* can help in investigating how sensitive the outcome of a model is to minor differences in probabilities [24]. A variety of sensitivity analyses exist, but in its simplest form a sensitivity analysis involves adjusting the probabilities of a variable and assessing its effect on the main probability of interest. This technique can aid probability elicitation, since it reveals how influential a variable is on the main variable of interest, as well as assist in the process of fine-tuning the parameters.

2.3.2 Model validity

The validity of a model defines how well it corresponds to the system it represents, regarding both the output and the mechanism that generates that output [39]. Testing the validity of a BN generally happens once it is completely specified, i.e. once its structure is complete and all probabilities are specified. The

procedures for testing the validity of a BN are usually quantitative in nature and regard the quality of the model's predictions (e.g. computing the Brier score [37] or accuracy of a model). In domains where little data is available, these measures will usually not suffice, because there simply is not enough data for a reliable evaluation.

The validity of models in data-poor domains is generally assessed by discussing the performance of the model with the domain expert(s) that contributed to its development [39]. Careful examination of the outcomes of the BN with different instantiations of evidence allows the domain expert to check whether the BN is a good representation of their view of the case.

Domain experts can disagree amongst each other on for example the modelling choices regarding the inclusion of variables, the values of the nodes, arc inclusion, arc direction, parameterisation and representation of (conditional) (in-)dependencies [56]. To allow for discussions about modelling decisions, it is important to thoroughly document any design choices or other steps of the modelling process (assumptions, motivations for building the model, sources of data). This however does not happen often in the literature. Efforts have been made to accommodate BN engineers and domain experts that are in discussion about modelling decisions for a particular BN, for example by Wieten and colleagues [56] using argumentation.

Part I

Guidance from the literature on manual BN graph construction

In this part of the thesis we summarise guidance provided by the literature on manually building the qualitative component of a BN model. Section 3 discusses in detail the construction process of the graph and expands on two approaches to graph construction for BNs from any domain. Section 4 elaborates on a method for BN graph construction for BNs regarding forensic DNA/biological traces. In other words, this is a method that is developed for designing highly domain-specific BNs, but it is also practical: as a stepwise construction method it specifies in detail which variables to include and how to connect them. This method is discussed because for our new guidelines we wish to achieve a similar level of practical applicability. Lastly in Section 5, some desirable properties for BN graphs specific to argumentation are discussed. Instead of stepwise guidelines, these properties could rather be considered a list of advantageous BN graph characteristics to keep in mind and work towards whilst establishing the structure. This list also functions as a checklist after completion of the graph. These properties form our main inspiration for proposing advantageous properties for BN graphs from any domain in Part II.

3 Manual graph construction – guidance from the literature

Section 2.3 included a brief insight into the general construction process of a BN. As in this thesis the focus lies on the qualitative component of BN models, it is useful to provide a more comprehensive synopsis of the steps involved in designing the BN graph (Sections 3.1 and 3.2). We also take a closer look at existing guidance the literature provides with respect to the development of the BN graph. Because the data-driven approach to graph construction is not applicable to data-poor domains, we focus exclusively on any guidelines on manual construction of the graph of a BN. The literature provides two general approaches, which are discussed in this section: Section 3.2.1 introduces a causal approach to building a BN graph, and Section 3.2.3 illustrates the use of idioms for designing a BN graph.

3.1 Identification of the variables and their values

The first steps in building a BN are identifying the relevant variables for the case to be modelled and specifying their values. Variables can be classified into different groups. These groups are merely useful for construction and organisational purposes: a BN is a mathematical model that does not recognise these classes and does not use them for performing probabilistic inference. Several types of nodes are distinguished by for example Jensen and Nielsen [24], Kjærulff and Madsen [26], and Korb and Nicholson [30] to capture the roles of each variable:

- *Hypothesis variables*: The main event(s) of interest and the one(s) of which the probabilities of its values are desired can be modelled in the hypothesis variable (also target, query, or output node [30], or problem variable [26]).
- *Information variables*: These represent variables which can be observed or measured and which possibly influence the hypothesis variable(s). Kjærulff and Madsen [26] further distinguish between:
 - *Symptom information variables*: these represent information that is available as a consequence of the main event of interest.
 - *Background information variables*: Aspects that can be viewed as causes for the main event(s) of interest or symptom variables captured in background variables.
- *Mediating variables* The third type of variables are mediating variables. These are unobservable variables, typically not causes, of which the probabilities are not of immediate interest. They are only introduced into the network for special purposes, such as to function as the middleman in order to correctly model independences, for computing inference more efficiently, and/or to enhance the graphical interpretation.

When the variables of importance are established, the values that variables can take should be specified. A variable’s values should not overlap in their potential to occur (mutual exclusivity) and should exhaust all possibilities (collective exhaustion). Preferably, values of a variables should not be mutually exclusive with one or more values of other variables, although this condition can be violated. While it is possible to include continuous variables in a BN, these are often discretised so their values represent non-overlapping intervals of the range of the original continuous variable [30].

Variables and values that are (eventually) insignificant to the case should not be included in the model, as they unnecessary complicate the network without being of any significant value in computing the probability of interest. It is possible that variables that were included at first are later ‘pruned’ from the structure in later iterations of the construction process.

3.2 Building the graph

Once relevant variables and their values are established, the next task is to draw arcs between variables in order to establish the network structure. There are two general approaches: drawing arcs based on the relations between variables, and establishing the network graph by means of idioms.

3.2.1 The notion of causality

It is important to remember that a BN graph merely represents the independence relation of a distribution at hand (see Section 2.2), and arcs do not carry any semantic connotation other than representing a possible correlation between two variables. However, when starting the construction process, thinking in terms of different types of relationships between variables can help in determining the appropriate network structure.

As a first heuristic, variables are often examined for any causal relations between them [24]. Causality is central in BNs [26] because it facilitates the process correctly capturing the independence relation amongst variables: a causal relationship typically corresponds to a conditional dependency. Moreover, another advantage of using the notion of causality is the possibility of intercausal interactions. The intuitive nature of causal relations makes a BN graph easier to explain as well as interpret. Lastly, a causal relationship between variables alleviates the task of probability estimation [26]. Estimating the probabilities for the occurrence of a particular effect assuming a known cause is easier than the other way around, as there can be other causes for a particular effect.

For these reasons, establishing arcs according to the notion of causality is generally suggested as a first heuristic by authors offering guidance to BN construction (such as [24, 26]). In case a variable has a direct causal effect on another, an arc is directed from cause to effect: specifically, an arc $V_i \rightarrow V_j$ is directed from V_i to V_j if V_i can be considered a direct cause of V_j . The relationship between two binary variables ‘Rain’ and ‘Wet grass’ for example can be

seen as causal: rain leads to the grass being wet. An arc is drawn from ‘Rain’ to ‘Wet grass’, as demonstrated in Figure 6.

Note that the structure in Figure 7 is not in line with directing arcs according to the notion of causality: wet grass does not result in rain. Although this structure is Markov equivalent to that in Figure 6 and mathematically there is no reason to choose one over the other, directing arcs as done in Figure 7 is strongly advised against in the literature (e.g. by Jensen and Nielsen [24]). Probabilities for ‘Rain’ have to be specified given ‘Wet grass’: wet grass can be caused by rain, but could also have been the result of another cause (a sprinkler for instance), which complicates the probability estimation.




Figure 6: A correct representation of a causal relation




Figure 7: Incorrectly representing a causal relation

Causal relations are not always obvious and are often subject to contentious debate, as Jensen and Nielsen [24] point out. Conflicting opinions can result in different BN graphs, even if these graphs exhibit the same independence relation (i.e. they are Markov equivalent). This is not uncommon as knowledge engineering is inherently subjective.

Kjærulff and Madsen [26] identify a basic structure that is typical for many BNs, based on a causal ordering between the different classes of variables established in Section 3.1. They argue that problem variables cause symptom and mediating variables (if any), and are in turn causally influenced by background variables. Background variables also can be seen as causes of symptom and mediating variables, and are therefore often root nodes in the network. If mediating variables are included, they often have a causal impact on symptom variables. This leads to the overall structure shown in Figure 8. Classifying the relevant variables from a case into the different categories and connecting them following the basic causal structure is one way of constructing the BN graph.

3.2.2 Other types of relationships

Aside from causal relationships, other types of relationships exist. These include association and temporal relations [30] and definitional relations. When knowledge of one variable provides information about another without any causal connotation, this can be seen as an associative relationship. An example is the association between age and literacy: age does have an effect on someone’s literacy in that statistically, the higher a person’s age, the more likely they are to possess the ability to read. However, this association should not be mistaken for causality: ageing does not *cause* someone to learn how to read. However

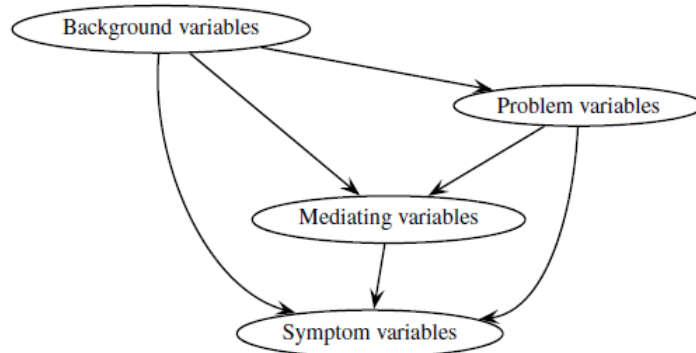


Figure 8: A basic structure for a BN, taken from [26, p. 155].

there is still an associative relationship between age and literacy, which can be captured by an arc.

Temporal relationships allow for a temporal ordering of variables. If a variable’s value changes before another, an arc can be drawn from the former to the latter so as to mimic the order in which events take place. Time plays a role in causal relations as well: in order for an effect to occur, its cause must have taken place earlier. Lastly, when a variable (an attribute) is defined by one or more variables (sub-attributes), they are in a definitional relationship.

The relationships between variables discussed in this section and Section 3.2.1 are merely helpful tools for BN graph construction. One has to keep in mind though that the arcs in a BN graph solely represent the (conditional) dependencies and independencies that hold in the modelled case. These d-separation properties should be checked after constructing the graph on whether they are correctly encoded to represent the views and knowledge of the domain expert [24]. Software tools exist that can help BN engineers with exploring these questions of whether (in)dependencies between variables in (initially) constructed BNs are correctly modelled, such as Matilda ([5] as described in [30]).

3.2.3 An idiom-based approach

Another approach to manual BN graph construction is based on the use of *idioms*, referring to fragments of BNs in which very generic types of reasoning are modelled. The idea is that complex inferential problems can be broken down into smaller components of generic reasoning patterns for which these idioms can be used as building blocks.

Idiom-based approaches are used in a variety of domains. In 1997, Laskey and Mahoney [32] introduced ‘network fragments’ and a framework for combining these in the field of military intelligence. In the same year Koller and Pfeffer [29] developed Object-oriented Bayesian networks (OOBNs), which allow a hierarchical way of organising a BN efficiently, akin to the framework of Laskey and Mahoney. Hepler [23] examined how an OOBN could be built for legal

cases. Quite similar to OOBNS, Van Gosliga and Van de Voorde [20] designed the Hypothesis Management Framework (HMF), which can be used for creating and extending BNs in a modular way.

Here we discuss idioms designed by Neil, Fenton and Nielson [35], who have developed idioms that are generally applicable to any domain. Fenton, Neil and Lagnado [16] and Lagnado, Fenton and Neil [17] also designed idioms specifically for the legal domain, these are instantiations of the general idioms and not considered further. Vlek and colleagues [51] also proposed idioms for modelling and combining scenarios in legal cases, which are not discussed in further detail here.

Neil, Fenton and Nielson identified five idioms in [35] for general manual BN graph construction. Each is discussed in more detail. The first idiom, the *cause-consequence* idiom from Figure 9, straightforwardly models a causal process from cause(s) to effect. The arc direction reflects this process. This concept is similar to the general heuristic of directing arcs in accordance with the notion of causality (as discussed in Section 3.2.1). Although the authors did not make use of idioms in constructing their BN, an example where this type of idiom can be identified is in the BN from De Ronde and colleagues [42], (see Figure 26 in Appendix A), where a suspect ‘S’ climbing the balcony (captured in the node ‘S climbed the balcony’) causes their fingerprints to be present on the balcony (node ‘Fingermarks S through climbing’).




Figure 9: Cause-consequence idiom (based on [35])




Figure 10: Definitional/synthesis idiom (copied from [35, p. 269])




Figure 11: Measurement idiom (copied from [35, p. 273])

The *definitional/synthesis* (Figure 10) idiom does not incorporate any causal notion, but rather represents a definitional relationship between the nodes (see Section 3.2.2). The arc is oriented in the direction in which a sub-attribute (or combination of sub-attributes) defines an attribute. An example is how someone’s paternal and maternal gene define their genotype, as in Dawid and colleagues’ simple paternity network [11] (see Figure 28 in Appendix A). Note that here again the authors did not make use of idioms in their network.

With the *measurement idiom* uncertainties regarding the accuracy of some observation can be modelled, whether this observation is done by ourselves, by others, or established by means of an instrument. In the measurement idiom in Figure 11, the ‘true value of attribute’ is the probability we are interested in. The ‘estimated value of attribute’ is instantiated: it is the result of a measurement done with a particular ‘estimation accuracy’ and acts as a “surrogate” for the true value, which is a latent variable and cannot be observed. An example can be found in Fenton and colleagues’ BN of the Simonshaven case ([18], see Figure 34 in Appendix A). They used an instantiation of the measurement idiom for modelling testimony evidence, where the ‘accuracy’ is the ‘credibility’ of the witness.




Figure 12: Induction idiom (copied from [35, p. 274])

When a sequence of results of a particular parameter are used to induce the result of another, perhaps slightly different, future parameter, the *induction idiom* (Figure 12) can be used. It is a way of modelling Bayesian inference, where context differences can be taken into consideration when the attribute or case in question is different from other historical cases (other children in this case). This idiom can be applied when inductive reasoning is the issue: suppose for example that we wish to know the effectiveness of a specific kind of treatment in a particular situation. Information is available on the application of this treatment in several previous cases, and a measure on the similarity or differences of these previous situations compared to the one at hand is available as well. With this idiom, one can induce the effectiveness of the treatment in the new situation from previous situations.

This idiom can also be used for instance to learn the probabilities for the estimation accuracy for a particular instrument from the measurement idiom. There are situations where no historical cases can be produced but the engineer is able to model the population parameter. In many cases an instance of the induction idiom in its full form is superfluous, as the population parameter can




Figure 13: Simplified version of the induction idiom, copied from [35, p. 275]

simply embody the distribution learnt from statistical data. In these cases a simplified version of the induction idiom can be used, as shown in Figure 13.




Figure 14: Reconciliation idiom (copied from [35, p. 276])

The last idiom models the reconciliation of a single node that is measured, defined, or produced by different methods (i.e. idioms or even BNs). In the reconciliation idiom in Figure 14, a node X results from two independent BNs (model A and model B). The reconciliation node is a binary node with values ‘true’ when the value of node X in model A is equal to that in model B, and ‘false’ otherwise.

The authors designed a flowchart (shown in Figure 15) to establish which idiom can be used in different situations. Along with this, the steps below describe the method they created for choosing and applying idioms.

1. Identify relevant variables
2. Identify relationships between these variables, resulting in subsets of variables
3. Examine these subsets in terms of the flowchart to determine the idiom which can represent the subsets.

Unfortunately the authors do not give any further information on how this flowchart was developed. Any motivation on why the questions are ordered this way is lacking.

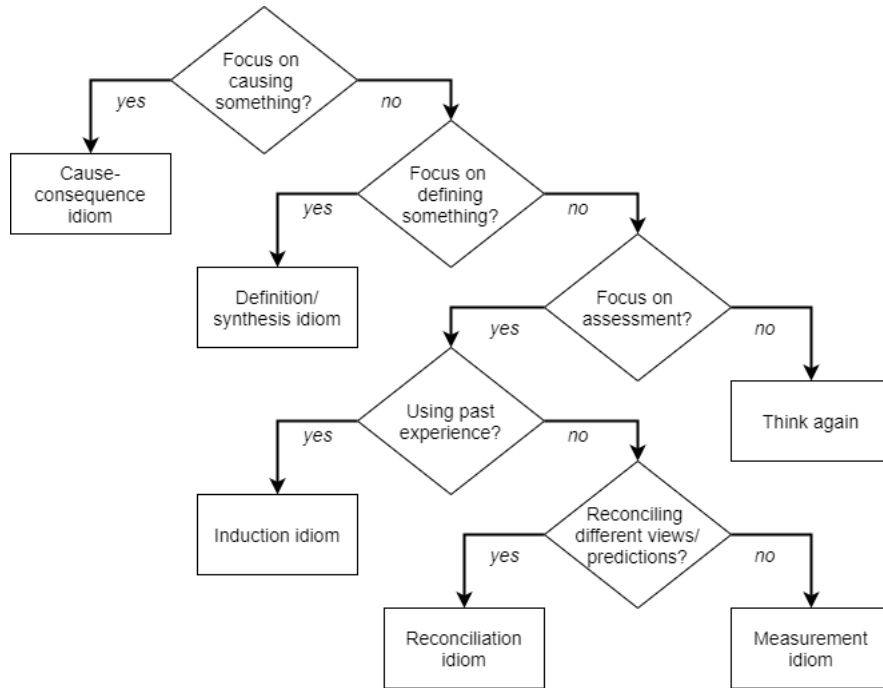


Figure 15: The flowchart for deciding what idiom to use (copied from [35, p. 278])

3.3 A note on model complexity

The qualitative component⁶ of a BN determines the complexity of the model. When building a BN graph, one has to strike a balance in finding the right level of complexity: a highly complex model in terms of many nodes, many parents and many values might lead to an excess in construction costs and infeasibility of performing probabilistic inference, but there must be just enough detail in the model to be able to perform meaningful computations whilst correctly representing the (in-)dependencies between the variables of the case. Adding nodes and arcs can increase the complexity of the model: the number of network parameters to be specified grows exponentially with the number of parents a node has. BN engineers should take this into account during the construction process. There are several ways to reduce the number of parents, if necessary.

One way is to remove “weak” arcs, which are those arcs that do not have a large influence on their children [30]. It is up to the domain expert and BN engineer to determine what constitutes a weak arc, but several methods exist to aid this decision (see for example [25] and [41]). Note that the influence of a variable on another by means of an arc can only be assessed after (initial) parameters have been specified.

⁶In this thesis we assume the qualitative component of a BN includes the values of variables.

Another technique is *divorcing*. In situations with many parent nodes for one child node, mediating nodes are introduced between a subset of all parents and the child node. Although this results in additional nodes in the structure and hence additional CPTs, it reduces the size of the CPT of the child node [24]. For example, a binary node with four binary parents requires $2 \cdot 2 \cdot 2 \cdot 2 \cdot 2 = 32$ parameters to be estimated, of which 16 are free. A binary node with only two binary parents requires only $2 \cdot 2 \cdot 2 = 8$, of which 4 are free parameters: dividing the four parents into two subsets and introducing binary mediating nodes between these and the child node results in less probabilities to be estimated, even though two new nodes are introduced. Figure 16 shows this effect.

In addition, in case of binary parent nodes, the states of the mediating nodes are often configurations of its parents. The parameters of these states are assigned ones if its corresponding parents are true, and zeroes otherwise. This does not complicate the process of probability estimation.

Caution has to be taken when applying the divorcing technique. Divorcing parent nodes from other parents is only possible when the impact of the divorcees on the effect node can be considered separately and independently from the impact of the non-divorced parents [15].

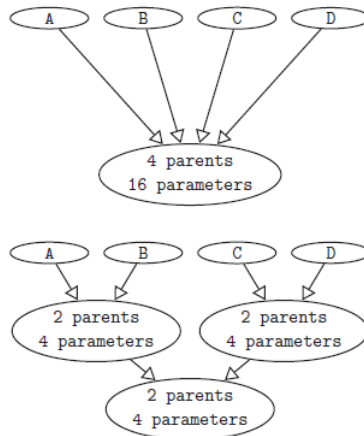


Figure 16: The introduction of a mediating variable reduces the number of parameters: assuming all variables are binary, the lowest node in the upper BN requires 16 free parameters to be estimated, the three lower nodes in the lower model combined require 12 due to the insertion of mediating variables (taken from [49, p. 94]).

4 BN construction in the forensic domain: A template by Taylor and colleagues

The guidance from the literature as discussed in the previous section aids the construction of BN structures in general. The guideline which is discussed in this section is highly domain-specific, but also very practical. It is useful to examine whether this specificity and practical applicability can be generalised for our own set of guidelines.

This section introduces a template by Taylor and colleagues [47] for building a BN graph by hand for forensic cases involving biological traces. Before discussing the template we first provide a short introduction on the necessary background information concerning forensic investigation and the use of BNs in this domain. The template is given in Section 4.3.

4.1 The evaluation of forensic traces

The task of a forensic practitioner is to examine the evidence (or findings) in light of a particular case with specific case circumstances. They can express their degree of confidence with a likelihood ratio (LR) for a set of *forensic hypotheses*⁷. It has become common practice that a pair of forensic hypotheses is defined where one forensic hypothesis represents the position of the prosecution in a case (denoted as H_p), while the other the one of the defence (H_d). The forensic practitioner evaluates the evidence, based on other information that has been provided, given these forensic hypotheses: how probable is it to find these results assuming forensic hypothesis H_p , respectively H_d is true? [45, 48].

Formulating a set of forensic hypotheses is not always easy or straightforward, which is why Cook and colleagues [10] specified a hierarchy of forensic hypotheses by classifying them into categories, or levels, to assist forensic practitioners in forensic hypotheses forming. Starting from the bottom of the hierarchy and going up, these are: (sub-)source, activity, and offence (or crime).

- *(Sub-)source level*: Forensic hypotheses on source level regard identifying the source of the evidence found, i.e. the donor of the trace, and determining the type of trace that is detected. An example of a pair of source level forensic hypotheses is:

H_p : Blood recovered from the victim's clothes came from the suspect

H_d : Blood recovered from the victim's clothes came from an unknown person

Even lower is the sub-source level, which deals with the source of the trace, but is concerned not with the cell type of the trace but only the DNA. This subtle distinction exists because it is possible to obtain DNA

⁷The standard term used in the literature is *propositions*. We however refrain from this term to avoid confusing these propositions with the general notion of propositions as used in BNs, namely as the values of variables [24].

profiles for traces of which it is uncertain from what body fluid or cell type they originate. The sub-source version of the previously stated forensic hypotheses is:

H_p : DNA recovered from the victim's clothes came from the suspect

H_d : DNA recovered from the victim's clothes came from an unknown person

- *Activity level*: The issue of interest for activity level forensic hypotheses is not only the who, but also the how: who deposited the evidence and what activity has resulted in the deposit of this evidence. Addressing these hypotheses involves not only assessing observations but also taking into account case-specific circumstances such as the transfer and persistence of the evidence. Below are examples of activity level forensic hypotheses.

H_p : The accused taped the victim

H_d : An unknown person taped the victim

- *Offence level*: At the offence (or crime) level, the forensic hypotheses encompass the question of whether or not a crime has been committed and who has committed it. This level closely resembles the activity level, with the important difference that these forensic hypotheses are often outside the domain of the forensic practitioner and up to the trier of fact (judge or jury) to evaluate. Example forensic hypotheses are:

H_p : The accused is the offender

H_d : Someone else is the offender

Thanks to technological advances over the past decades, identifying the source of a forensic trace, such as the donor of DNA, has become very reliable. Therefore the source of a trace is rarely challenged these days, but rather the manner of deposition of the trace, i.e. the activity that led to the trace ending up where it did, is put to question [27].

4.2 The use of BNs in the forensic domain

Since forensic hypotheses are a central component of the interpretation of forensic findings, BNs in the domain of forensic science and law are often constructed around a set of forensic hypotheses. A BN is built that propagates all evidence and returns the LR, often in a separate forensic hypotheses node, which can be used to communicate an informed decision on the interpretation of evidence under either forensic hypothesis. An advantage of BNs is that they can be updated easily when provided with new information. During construction this ability is useful, because BN engineers can monitor more easily how adding evidence influences this LR.

Where the focus in the past has mostly been on source-level BNs, i.e. BNs built for the evaluation of evidence under a set of forensic hypotheses that can

be classified as source level (e.g. [11, 34, 50], recently it has shifted towards activity level BNs (e.g. [42, 44, 47, 55]).

The hierarchy of forensic hypotheses is of relevance to BN construction in the forensic domain because the forensic hypotheses make up the ultimate hypotheses in the case to be modelled. Parallels can be drawn between BNs with forensic hypotheses from the same level. For example, variables that are often included in source-level BNs are number of contributors to a DNA trace (e.g. [34], [59]) and test results, e.g. for body fluid types ([36], [46], [57]).

For activity-level BNs, variables that are often included are transfer and persistence of biological material ([28, 44, 55, 57]), contamination of the sample ([4, 55]), trace characteristics (e.g. sampling location ([2, 55]), visual aspects ([22]), location ([42])), and sample donor characteristics ([22, 42]).

4.3 A template for activity level BN construction for biological traces

With the increase in activity-level evaluation of evidence in forensic investigation and the rise in the use of BNs for this type of cases, Taylor and colleagues [47] recognised both the relevance and difficulty of BN construction. To tackle this, they came up with a stepwise method of constructing BNs for forensic biology cases with activity level forensic hypotheses:

1. *Define main hypothesis node*: determine the forensic hypotheses and include them as the values of a single node, which is the main (hypothesis) node of the model.
2. *Define activity nodes*: identify the activities in the case and draw a node for each of these. Make them children of the hypothesis node. Taylor and colleagues advocate the inclusion of non-disputed activities (which are true under both H_p and H_d). These can be seen as redundant nodes, unnecessarily complicating the structure, but this way the model can clearly represent the influence of undisputed activities on findings nodes.
3. *Group similar findings*: if possible, combine (sets of) findings of similar samples such that they can be considered a single item. The samples should be similar in the sense of originating from closely related areas (e.g. a set of tape lifts taken from different areas of the same item).
4. *Define findings node(s)*: draw a single node for each group of similar findings. These nodes are eventually instantiated in the model. No arcs are drawn yet in this stage to connect them to the other nodes.
5. *Define transfer and persistence nodes*: these nodes describe the mechanisms of how the activity resulted in the findings and are placed between the activity and findings nodes. Connect the nodes while taking the following into account:
 - There can be more than one activity contributing to a single result.

- There can be more than one transfer and/or persistence node in a chain from activity to findings node.
 - The order of activities and transfer is potentially of importance.
 - Mediating nodes can be included to increase the comprehensibility and reduce the complexity.
6. *Define root node(s)*: these nodes represent variables that do not relate to any activity but are relevant (as parents) to transfer and/or persistence nodes (e.g. background levels of fluids, contamination).
7. *Check for absolute support within the BN*: this also relates to probability estimation; avoid obtaining a BN where one piece of evidence can, upon instantiation, result in a single forensic hypothesis being true while the other is false. To avoid this, Taylor and colleagues advise not using probabilities 1 and 0 in the CPTs of the transfer and persistence nodes (an exception is made for the values of mediating nodes). Additionally, the BN should be checked on whether all findings are observable under each forensic hypothesis. This means that there should be multiple chains leading to a findings node, at least one corresponding to the events of one forensic hypothesis and one to those of the other forensic hypothesis. In practice there are typically two competing causes for a findings node, these causes take the form of either two competing activity nodes, or a competing activity and root node. If there is only one chain to a findings node, the BN is at risk of having a probability of 1 for one of the forensic hypotheses once the evidence is instantiated.

Taylor and colleagues applied their method to an example case and the resulting BN is shown in Figure 25 in Appendix A. De Ronde and colleagues [42] also use this method, see Figure 26 in Appendix A.

5 Advantageous properties for BN graphs constructed from arguments

Previous sections discussed BN graph construction guidelines from the literature, with general approaches in Section 3 and a domain-specific method in Section 4. In this section we discuss another form of guidance to BN construction the literature provides, in the form of a list of advantageous properties⁸ for BN graphs from Timmer [49]. Timmer proposed this set of properties for BNs for argumentation schemes. These properties do not serve as a stepwise method for BN graph construction, but rather as a list of BN characteristics that should or can be achieved during construction and as a checklist for evaluating the structure after (initial) construction.

In order to understand Timmer’s properties, we briefly introduce some concepts from argumentation. In argumentation, an argument consists of one or more premises and a conclusion. A premise provides support or a reason for drawing a particular conclusion. Naturally, a conclusion follows from one or more premises. Timmer specifically addresses BN construction of argumentation schemes⁹, which are formal structures that capture a general pattern of reasoning, such as *modus ponens*, or *argument from position-to-know* [53]. These schemes are often accompanied by a set of ‘critical questions’, which address weak points in the argumentation scheme where it can be attacked, or which question the application of the argumentation scheme [53].

The advantageous properties, copied from p. 96, are listed here with a short explanation.

1. *Critical questions are explicitly modelled by a node*: this property supports the inclusion of an explicit node for each critical question that is of relevance for an argument in a particular situation.
2. *Critical questions can explain away the hypothesis (via a [converging connection])*: this structure allows for the possibility of intercausal interactions (specifically, explaining away) between the hypothesis, i.e. the conclusion of the argument in the argumentation scheme, and the critical questions. For this to happen, the hypothesis and critical questions require a common child node (the premise) which is either instantiated or has an instantiated descendant.
3. *The number of free model parameters is as small as possible*: this property addresses the need for reducing the computational complexity and keeping the task of probability estimation doable. In order to reduce the number

⁸Timmer uses the word ‘criteria’, even though these are not meant as strict norms a BN has to adhere to. For this reason we use the word ‘properties’.

⁹Timmer’s properties only apply to embedding critical questions in a BN for a selection of argumentation schemes, namely evidential argumentation schemes. This method is not straightforwardly applicable to predictive argumentation schemes or other types. We will not discuss this distinction further here, but refer the reader instead to Timmer [49].

of model parameters, one should strive to keep the number of parents to a minimum, as well as limit the number of possible values for a variable.

4. *No redundant [arcs] are included:* additional arcs complicate the model and increase the number of parameters to estimate. Furthermore, this property prevents unwanted interactions between variables of which it is known they do not interact or of which is uncertain whether they do.
5. *Premises and conclusions are explicitly modelled by a node:* it is advisable, if not essential, when modelling arguments in a BN to include nodes for the premises and conclusions.
6. *There can be an active chain from the hypothesis to the evidence given the variables that can potentially be observed:* following this property allows for the evidence to influence the hypothesis. An active chain is necessary, unless it is desired that the effect of evidence is undone given other variables that are observed (which block the chain upon instantiation).
7. *When applicable, [arcs] follow the direction of temporal/causal precedence:* this property captures directing arcs according to the notion of causality (discussed in Section 3.2.1) or in a temporal manner, from events that occur earlier to later ones.

The first four properties are formulated specifically for those BNs that model (evidential) argumentation schemes and incorporate critical questions. Properties 5 through 7 relate to the general construction of BNs from information expressed as arguments or argument schemes, i.e. whether they include critical questions or not.

Part II

Formulating new guidelines and comparison measures

In the previous sections we have provided an overview of the construction process for the qualitative aspect of a BN. After a brief introduction of the relevant concepts of BN theory and its general building process in Section 2.2, we concentrated on the manual construction of the BN graph. Several guidelines from the literature were summarised: Section 3 expanded mainly on causal arc direction and an idiom-based approach to graph construction. A highly domain-specific method by Taylor and colleagues [47] is described in Section 4. Finally, Section 5 provided not so much a method, but a list of desirable properties by Timmer [49], again for domain-specific BN graphs.

In Part II we aim to formulate a set of practical guidelines for manual BN graph construction for cases from any domain, as well as a list of measures to compare BN structures by. Because of the practical nature of Taylor and colleagues' template, it could serve as a good inspiration for our own guidelines. However, first we have to examine whether this template is even remotely applicable to problems from other domains. To do so, we have to find a way to assess the BN graphs resulting from this template. Timmer proposes a list of properties. These are again specific, but in contrast to the template, specific to argumentation instead of to a particular domain. Nevertheless, they could be informative for the properties any BN should possess.

In this part we first describe the process of collecting raw data on handcrafted BN structures from observations from the literature and information obtained through an interview with a BN engineer and domain expert (Section 6). We use these observations from the field throughout Parts II and III to base our conclusions on. We further aim to generalise Timmer's properties in such a way that any manually constructed BN graph can be analysed on these properties (Section 7). In judging how to generalise these properties we use the general graph construction approaches described in Section 3 and the obtained BN graph knowledge (Section 6) for this purpose. The resulting new properties form our main inspiration for our comparison measures. In addition, these properties describe desirable characteristics of BN graphs. As such they can be informative in designing our guidelines, as we aspire that the BNs resulting from our guidelines possess these properties. In Section 8 we explore whether the BN structures resulting from Taylor and colleagues' template encompass these new properties to examine whether this template can help in the formulation of the guidelines. Finally in Section 9 we propose our set of guidelines, and in Section 10 our comparison measures.

6 Manual BN graph construction - Obtaining observations from the field

In order to get acquainted with the modelling conventions and motivations behind modelling choices in practice, we made the effort of studying what actual handcrafted BN graphs from the field look like and how these are constructed. We did so in two ways.

Firstly, we gathered information from existing BN graphs that were constructed by hand for problems from two data-poor domains, namely the forensic and legal domains. From 32 articles ranging from 2002 to 2019, a database of 69 BNs was created (see Table 4 in Appendix B for a complete overview). All BN graphs are hand-crafted. We examined these structures on several aspects:

- Numerical characteristics: the number of nodes, number of arcs, number of head-to-head nodes, maximum number of incoming arcs (highest in-degree), number of values
- Approaches to graph construction and modelling choices
- Motivations for constructing the BN
- Level in hierarchy for the forensic hypotheses, if any (see Section 4.1)
- Variable relationships and arc direction
- (Conditional) dependencies and independencies represented by the graph
- Origin, i.e. from which institute the authors originate

In addition, a selection¹⁰ of the BNs was examined on adhering to an adapted version of the properties by Timmer (see Section 5).

This in-depth investigation of actual BNs formed a solid basis for our knowledge of what BN graphs look like in practice. The graphical aspects we studied provided a nuanced image of the conventions and considerations in BN graph construction in these domains.

Secondly, a semi-structured interview was conducted with a forensic scientist with DNA expertise at Netherlands Forensic Institute (NFI), where manual construction of BNs is widely used. The interviewee can be considered as both an expert in the forensic domain as a BN engineer. The purpose of this interview was to gain insight into the methods employed by this department at the NFI in construction of BNs, specifically BN graphs. A concise outline of the interview is given in Section B in Appendix B.

These two approaches together serve as ‘raw data’, which is used in subsequent parts of this thesis. The insights gained from collecting and assessing this data is used to base our conclusions on when examining the theoretical methods.

¹⁰Due to time constraints, we could not examine the database in its totality on this aspect.

7 A new set of advantageous properties for BNs from any domain

This section explores whether Timmer’s properties (Section 5) can be generalised such that they are applicable for BNs from any domain. We make use of the general graph construction approaches described in Section 3, as well as additional literature and the practical knowledge from the field (see Section 6) in deciding if and how to adapt Timmer’s original properties. By individually assessing each of Timmer’s properties we formulate a new set of properties. Afterwards we also inspect whether it is achievable for a BN graph to possess all new properties. In extending the applicability of Timmer’s original properties to a new set of non domain-specific properties we find inspiration for our own guidelines and comparison measures.

Property 1: Critical questions are explicitly modelled by a node

According to Timmer, a reason for modelling critical questions as explicit nodes is to facilitate explaining-away, which is dealt with in further detail when discussing Property 2. Another reason is that other evidence can be connected to a critical question node that pertains to that question only.

To illustrate this, suppose that an expert E asserts that a particular hypothesised event A (which falls under E ’s domain) is true. From this, one might conclude that A is true. This is an informal explanation of the Argument from Expert Opinion scheme, from Walton [53]. A critical question associated with this argumentation scheme concerns E ’s expertise: ‘How credible is E as an expert source?’, questioning whether A can be validly concluded from E testifying to A . In case there is evidence of E being partial for example, E ’s credibility is put into question.

Figures 17 and 18 model this argumentation scheme, where ‘H’ is the hypothesised event A , ‘CQ’ is the critical question, ‘Exp’ is the expert’s testimony to A , and ‘Bias’ represents the evidence of a bias. In Figure 17 the hypothesised event and critical question are modelled together in a single node. The evidence ‘Bias’ is only relevant to the critical question, but since probabilities have to be estimated for every configuration of values of the parent node(s), the probability of this evidence given the hypothesised event has to be estimated. In Figure 18 the hypothesised cause and critical question are modelled separately, such that ‘Bias’ is only connected to ‘CQ’. Meaningless probability estimations are avoided in this way.

In order to generalise this property, we take a look at the role that critical questions play in this type of argumentation schemes: they can be viewed as causes that compete with the hypothesised cause. Evidence resulting from the hypothesised cause can also emerge as a result from a critical question. For the argumentation scheme in the example above, the conclusion (A in this case) can be seen as the hypothesised cause, and the premise (E ’s testimony to A) as the




Figure 17: Collapsing the hypothesised cause and critical question into a single node




Figure 18: Separate modelling of the critical question as a single node

observed effect. However if E turns out to be not a credible source, this could also ‘cause’ E to assert A . The critical question can be viewed as an alternative cause to the effect.

Timmer models these “alternative causes” as explicit nodes. Causes do not necessarily have to be critical questions. We consider causes more broadly.

As mentioned in Section 3.2.1, BN graphs are often built with the arc direction in the first place according to the notion of causality. Neil and colleagues idioms include an idiom to model cause and effect (the cause-consequence idiom, Section 3.2.3). The way in which causes are modelled is important for BNs from any domain. In this light, we extend property 1 such that causes are modelled by separate nodes.

New property 1. *Causes are explicitly modelled by separate nodes.*

The reason to model causes in separate nodes is the same as Timmer describes for his first property: collapsing competing causes into a single node poses a problem. Compare Figures 19 and 20, which model the same case. In this case, adapted from Fenton, Neil and Lagnado [14], two alternative causes $C1$ and $C2$ lead to a common effect E . In addition, there is a variable $V1$ that only leads to $C1$, which in turn only results in evidence $E1$.




Figure 19: Modelling competing causes as a single node




Figure 20: Modelling competing causes as separate nodes

The main problem with modelling the causes as distinct values of a single node, as done in Figure 19, has to do with the causes and evidence being part of a distinct causal chains. By collapsing the alternative causes into a single node, we are forcing the causes to be part of each other’s chain. As a consequence, meaningless probabilities have to be estimated, which we will explain in more detail.

Suppose the states of the *Causes* node in Figure 19 are $\{c1, c2\}$. In addition to the fact that the possibility of neither (or both) causes occurring is ruled out, this way of modelling is problematic for correctly filling out the CPTs. The reason is that the causes are part of each other’s chain, specifically the chain between *E1* and *V1* via *Causes*. *V1* is relevant for cause *C1* and therefore the probabilities $\Pr(c1 | v1)$ and $\Pr(c1 | \neg v1)$ will (likely) differ. The probabilities of *C2* given *v1* or $\neg v1$ should be the same, i.e. $\Pr(c2 | v1) = \Pr(c2 | \neg v1) = \Pr(c2)$, because the occurrence of *C2* should not be influenced by *V1*. Not only are we forced to consider *V1* when estimating the probabilities of *C2*, also as can be seen in Table 2, $\Pr(c2 | v1) \neq \Pr(c2 | \neg v1)$ because the probabilities for each (conditional) distribution in a CPT must sum to 1.

Similarly, estimating the probabilities for the CPT for *E1* involves the pointless consideration of *C2*, making these parameters meaningless (see the question marks in Table 3). Lastly, specifying the values of the *Causes* node this way rules out the possibility of the presence of both causes, as well as their joined absence, which can be relevant for *E*. Even adding these as values does not resolve the issues mentioned above.

<i>V1</i> →		<i>v1</i>	$\neg v1$
<i>Causes</i>	<i>c1</i>	x	y
	<i>c2</i>	$1 - x$	$1 - y$

Table 2: CPT for the *Causes* node from Figure 19, with values $\{c1, c2\}$

<i>Causes</i> →		<i>c1</i>	$\neg c2$
<i>E1</i>	<i>e1</i>	x	?
	$\neg e1$	$1 - x$?

Table 3: CPT for node *E1* from Figure 19, with values $\{c1, c2\}$ for the node *Causes*

Explicitly modelling alternative causes as separate nodes, as done in Figure 20, avoids irrelevant state combinations. *V1* and *E1* are now only part of *C1*’s chain, such that no meaningless probabilities have to be estimated.

Another valuable feature of modelling causes individually is that it is often beneficial to the representational aspect the BN graph provides. Separate nodes often allow for a quicker overview of which variables have been considered for a node than when they are modelled as distinct values of a single variable. Of course there is a limit to this representational profit: too many causes might overcomplicate the network visually.

A difficulty arises when two causes are mutually exclusive: preferably, the values of a nodes must not be mutually exclusive with those of another variable. This property calls for separate nodes for the causes, but mutual exclusivity is

not enforced in this way. To illustrate, suppose E in the example above stands for ‘Missing lunch’ from the fridge at work. E and E represent ‘Someone ate it’ and ‘Someone threw it out’. E and E are mutually exclusive: if one is true, the other cannot also be true¹¹. Property 1 however calls for separate modelling as done in Figure 20, which leads to the forced estimation of the impossible probability of $\Pr(\text{missing lunch} \mid \text{someone ate it} \wedge \text{someone threw it out})$ for the CPT of E . In other words, the event that both are true is not excluded with this way of modelling.

To solve this issue, several options have been proposed, for example by Korb and Nicholson [30] and Jensen and Nielsen [24]. The collapsing of the causes into a single node is, as we have just shown, not preferred. An arbitrarily directed arc could be drawn between the two causes to introduce a dependency, such that it does not matter what numbers are specified in the CPT for E . There is no natural arc direction, and again the problem of non-distinct chains emerges: in case the arc is drawn from E to E , for any other variables leading to E , the probabilities of E must be estimated with E in mind.

Fenton, Neil and Lagnado [14] propose a satisfying solution: adding a deterministic constraint node as a mediating node, with its parents as values and an additional value ‘NA’, not applicable. This last value ensures mutual exclusivity by being equal to 1 for impossible state combinations¹². Figure 21 illustrates this constraint node is introduced between the competing causes and effect node from the previous example. We find this is the most elegant way of enforcing mutual exclusivity whilst keeping cause nodes separate. Additionally, it works for any number of competing causes. This approach is a useful feature for our guidelines.




Figure 21: Enforcing mutual exclusivity between mutually exclusive causes, where ‘Con’ is a constraint node, copied from Fenton, Neil and Lagnado [14]

¹¹Here we make the assumption that the lunch was tasty, such that someone starting it will have finished eating it and will not have thrown it out halfway.

¹²For details on how to specify the probabilities we refer the reader to Fenton, Neil and Lagnado [14].

Property 2: Critical questions can explain away the hypothesis (via a converging connection)

Thanks to property 1, critical questions are explicitly modelled by separate nodes. The event (A in the example in the previous section) is the hypothesis. The hypothesis and critical question(s) together can form a converging connection, with the common effect as the head-to-head node. This allows for the possibility¹³ of the explaining away effect.

The explaining-away effect is of course not only useful in the domain of argumentation. Allowing for the possibility of it is also beneficial for cases from other domains. Building on the previous paragraph on property 1, we could then straightforwardly adjust this property so that alternative causes can explain each other away, via a converging connection. In generalising this property however we can aim broader: any type of intercausal interaction should be captured by the BN graph. We can therefore adapt this property as follows:

New property 2. *The network structure can allow for intercausal interaction.*

Property 3: The number of free model parameters is as small as possible

In the interest of reducing the computational complexity and keeping the task of probability estimation doable, property 3 should be pursued for any BN. The number of (free) parameters to be estimated increases exponentially with the number of parents for a node (see Section 3.3). Keeping this number of parents as small as possible reduces the number of parameters to be specified and with that increases the computational feasibility of the model.

Timmer specifically aims to reduce the number of *free* parameters, as these are the ones that have to be explicitly specified (in contrast to the bound ones, whose numbers result from the specification of its free counterparts). We keep the property as is:

New property 3. *The number of free model parameters is as small as possible.*

Following this property amounts to both minding the number of states for a node, as well as minding the number of parents. Additionally, keeping the CPT sizes acceptable lowers the number of probabilities to estimate. Section 3.3 has already dealt with techniques to confine the number of parents of a node to a minimum.

This property is difficult to follow in case mediating nodes are included in a BN. These often have all possible parent configurations as values, which can lead to a large number of states. Probability estimation and model complexity however are often not an issue, since mediating nodes are nearly always deterministic. Since mediating nodes therefore help in reducing the model complexity and this is the goal of this property, we propose that the total number

¹³Whether explaining away actually occurs of course depends on the specification of the parameters.

of free parameters does not have to include the parameters from the CPTs of deterministic mediating nodes.

Property 4: No redundant arcs are included

Like property 3, property 4 also strives to keep the computational complexity under control. Arcs bring along parenthood and with that larger CPTs. Furthermore, this property prevents unwanted interactions between variables of which it is known they do not interact or of which is uncertain whether they do. Aiming to adhere to this property will stimulate the BN engineer to think critically of the influences of variables on others.

Determining whether an arc is redundant involves considering the influence between the connected variables and the representation. Weak arcs (see Section 3.3) can be removed to strive to follow this property. BN engineers should keep in mind though that removing an arc removes the direct dependency between two variables, but does not exclude the possibility of an indirect effect through another active chain.

Property 4 is directly applicable to BN graphs from other domains. We therefore keep it as is.

Property 5: Premises and conclusions are explicitly modelled by a node

‘Premises’ and ‘conclusions’ are well-defined terms in argumentation. It proves difficult to directly generalise this property, as it is complex to find a general term for any concepts from other domains that are similar to ‘premise’ and ‘conclusion’.

One possibility is to view premises as observed evidence and conclusions as hypotheses. This is not straightforwardly so, as for example in some types of reasoning premises are not necessarily observed. Additionally, there can be intermediate conclusions which can serve as a premise for a later conclusion, although this was not the case for the specific argumentation schemes that were the subject of Timmer’s research. We therefore refrain from generalising this property, but allow it to inspire us to form our own.

For a BN to function properly, nodes should be included for the main hypothesis of interest, as well as for observed evidence. Any other nodes connected to these are not strictly necessary for the BN to function. We argue every hypothesis and every piece of information that is or can practically be observed in the case should be modelled as a separate node, and not as values of single nodes, in similar fashion to property 1. This increases the value of a BN graph as a representational tool by providing a clear overview of the variables considered for the model. The argument of distinct chains (see property 1) also applies here. We choose to encase this in a property:

New property 5.

5.1 Hypotheses are explicitly modelled by separate nodes

5.2 Evidence is explicitly modelled by separate nodes

In many domains the main hypothesis or hypotheses of interest can be viewed as cause(s) for the observed evidence. Modelling these as separate nodes is already covered in property 1. This is not always the case: in some BNs, the main hypothesis of interest is more a result of a series of other (possibly observed) causes, for example when examining the status of a phenomenon or determining the utility based on (a combination of) different variables. We therefore choose to draw a property 5.1 explicitly for hypotheses.

Property 6: There can be an active chain from hypothesis to evidence given the variables that can potentially be observed

This property should be respected, indeed if there is no possibility of an active chain between hypothesis and evidence, the evidence does not influence the hypothesis and should be left out. Theoretically any node in the model can be instantiated, but in practice only nodes which can be observed in the case will be. These should not block chains between other observed evidence and the hypothesis node(s).

Only if it is absolutely desired that there be an independence between observed evidence and hypothesis upon instantiation of other nodes in the chain, can one allow for blocked chains. There cannot be any uncertainty regarding the observed variable that blocks the chain, since it will fully prohibit any influence between the hypothesis node(s) and the evidence node(s) on that chain.

As this goes for any BN graph, we keep this property in its original form.

Property 7: When applicable, arcs follow the direction of temporal/causal precedence

This practice stimulates a clear understanding and interpretation of the BN and additionally facilitates the construction of the network, as well as probability estimation. This property reflects the general tendency of causal arc orientation as a first heuristic of drawing arcs as favoured by authors in the literature (e.g. [24], see Section 3.2.1).

As mentioned in Section 3.2.2, variables can be not only in a causal or temporal relationship, but also in a definitional relationship. This last relation mirrors the way the arcs are directed in the definitional/synthesis idiom (See Section 3.2.3, Figure 10), namely from sub-category to category or sub-attribute to attribute. We can therefore choose to extend this property as follows:

New property 7. *When applicable, arcs follow the direction of temporal, causal or definitional direction.*

A note on the internal coherence of the list of properties

The properties, both in their original form and in their adapted form, have disagreeing goals. Trying to adhere to all properties discussed above is problematic. Those properties supporting representational value of the graph clash with the ones aiming to minimise the computational complexity. Specifically, property 3 is in conflict with aiming to follow properties 1, 5.1 and 5.2: separate modelling of variables results in more parents, with which the size of the CPT of the child node grows exponentially, leading to a higher number of (free) parameters to estimate. Property 1 also collides with property 4, because if causes can be modelled as a single node, one could argue that modelling them separately leads to redundant arcs.

We can conclude from this that, when building a BN graph, there is a trade-off between computational complexity and visual representation of the modelled case. Both are important, but in some domains one might weigh more than the other. In the forensic domain, where BNs are used as a tool to communicate how a LR is established, representation might carry more value than in domains where the modelled cases include a myriad of nodes and arcs. In such domains computational feasibility might matter more, such as for some cases from the medical domain with many symptoms and influential variables.

7.1 Conclusion

In the previous paragraphs we discussed how Timmer’s list of advantageous properties can be adapted to formulate new properties for BNs from other domains. We kept properties 3, 4 and 6 in their original form. To summarise, adding a ‘P’ to distinguish the new properties from Timmer’s properties:

P1 Causes are explicitly modelled by separate nodes.

P2 The network structure can allow for intercausal interactions.

P3 The number of free model parameters is as small as possible.

P4 No redundant arcs are included.

P5.1 Hypotheses are explicitly modelled by separate nodes.

P5.2 Evidence is explicitly modelled by separate nodes.

P6 There can be an active chain from hypothesis to evidence given the variables that can potentially be observed.

P7 When applicable, arcs follow the direction of temporal, causal or definitional precedence.

We feel these properties reflect the important issues to keep in mind during construction: representational value of the BN structure, computational complexity, facilitation of parameter estimation, and the basic requirements for BNs to function.

8 Assessment of the template by Taylor and colleagues by means of the new properties

This section takes a closer look at Taylor and colleagues' method [47] (Section 4) to assess whether following the steps results in a BN graph that incorporates the new properties as adapted from Timmer [49] (see the final list in Section 7.1). The template is designed specifically for BN graph construction for forensic biological traces given activity-level forensic hypotheses. The adapted properties should be applicable to BNs from any domain and therefore should also pertain to the BN graphs resulting from Taylor and colleagues' template. By assessing the steps in Taylor and colleagues' method to see how they relate to the new properties, we can draw inspiration for the guidelines that we will formulate in Section 9. In addition, any properties not incorporated in the BNs from this method demonstrate gaps in the method, which we can learn from when drafting our guidelines.

We go through the steps as prescribed by Taylor and colleagues, denoted as for example 'Ta1' for step 1 of the method.

Ta1: Define main hypothesis node

By formulating a main hypothesis node, the BN will have a single node which models the activity-level forensic hypotheses. These forensic hypotheses are modelled as distinct values of a single node, which is not in line with property P5.1. According to this property, separate nodes should be included for every hypothesis.

Nonetheless, as discussed in Section 4.1, forensic practitioners typically use a BN to formulate likelihood ratios (LRs), which are presented to the trier of fact. These LR's can be directly read from the hypothesis node. Therefore, this specific sub-domain of the forensic domain might benefit from combining the forensic hypotheses in a single node.

Ta2: Define activity node(s)

The activities from the case can be considered different explanations of the evidence, which can be seen as alternative causes for that evidence. Since the method calls for drawing a node for each activity, this is in line with property P1.

By also including undisputed activities as nodes, the method does not necessarily follow properties P3 and P4: a node that has equal probabilities under both forensic hypotheses is essentially redundant and can be omitted (and consequently redundant arcs are removed and the number of parents for other nodes decreases). The authors state that this is an option: omit the arcs from these activity nodes to the hypothesis node or omit these nodes altogether and adjust the probabilities in the other nodes accordingly. However, the increase in comprehensibility and explainability of the BN is worth the loss of computational

efficiency, according to the authors.

Ta3: Group similar findings & Ta4: Define findings node(s)

Not including every single piece of evidence as a node in the model but instead grouping similar samples together is an improvement with respect to the computational complexity. This is in line with properties P3 and P4.

Property P5.2 calls for explicit modelling of evidence: every piece of evidence should be included in a separate node. If, however, different pieces of evidence originate from the same area and share a similar influence on other variables, modelling all these pieces individually needlessly complicates the model. The graph then includes a myriad of nodes representing these pieces of evidence and complex dependencies between them which are not particularly informative and difficult to estimate. Treating them as a single piece of evidence is therefore acceptable, and does not completely contradict P5.2.

Ta5: Define transfer & persistence node(s)

In this step the activity nodes are connected to the findings nodes through the transfer and persistence nodes. If more than one activity is connected in this way to a single result in a findings node (which is eventually instantiated), this allows for an indirect explaining away effect between the competing activities. This is in line with property P2.

Mediating nodes can be added to summarise the transfer and persistence nodes. This can help in reducing the number of parents for a single node, and with that the number of parameters (property P3). The BN in Figure 25 in Appendix A contains a mediating variable such that the findings node ‘family YSTR profile [...]’ has two parents instead of three.

Lastly, none of the activity or transfer and persistence nodes are instantiated. Because of the way the arcs are drawn, this results in there being at least one chain between the forensic hypotheses and each evidence node, forming a series of serial connections that is not blocked. There is an active chain from hypothesis to evidence, so this conforms to property P6.

Ta6: Define root node(s)

Depending on their definition, the role of root nodes is often as a cause for a piece of evidence that competes with the activity node(s) connected to that evidence. By including root nodes as separate nodes, this step is in line with property P1. Additionally by making them a parent of the findings node, explaining away can take place (property P2).

Again Taylor and colleagues suggest that the root nodes *can* be omitted whilst adjusting the probabilities in the findings nodes. Doing so would be in pursuit of properties P3 and P4, but again the authors vouch for including them to increase the representational value and comprehensibility.

Ta7: Check for absolute support in BN

In this step the BN is checked whether all results are observable under either forensic hypothesis. For this to happen, there must be first of all no chains between evidence and hypothesis that are blocked by variables that will be instantiated given the case (i.e. evidence). Property P6 dictates this feature.

Furthermore, this step explicitly calls for the inclusion of multiple causes for a findings node. Generally there are two competing causes, each supporting one of the forensic hypotheses. As such, property P2 can be incorporated in this step, although there is no explicit mention of explaining-away.

Taylor’s method in general

As a result of starting at step 1 and ending at 7, the general tendency of directing the arcs is from hypothesis to activity nodes, to transfer and persistence and eventually findings nodes. In other words, arcs are directed in a causal or temporal manner (property P7). Drawing arcs from root nodes (as alternative causes) to findings also follows the notion of causality. Finally, introducing a mediating node as a child from transfer and persistence nodes mirrors a definitional arc direction: the transfer and persistence nodes define the mediating node.

Apart from the possibility of including mediating variables, there is no explicit effort to adhere to property P3, i.e. to keep the number of parameters as small as possible by minding the number of values for a node or the number of parents. Ways to do so are described in Section 3.3. It is possible to conclude that this method does not aim to achieve the best conceivable computational efficiency.

The evaluation of the graph is subtle. The resulting graph is not evaluated on the basic requirements for a BN graph, i.e. no cycles and a correct representation of the independence relation amongst the variables. We note that the way these steps are designed, it follows that cycles will generally not occur, and in most cases the (in-)dependences will be correct. However, if in a particular case there are more complex dependences between variables, the method does not provide any guidance. Additionally, the method does not offer room for iteration: options to return to earlier steps are not explicit, although the BN graph construction is usually iterative [13].

BNs with activity-level forensic hypotheses are aimed at answering questions regarding the source and manner of deposition of traces. However, this method enforces that these questions are answered in the form of activities, which therefore have to be clearly formulated. For complex problems (e.g. the question of whether or not a piece of tape contains a DNA sample from a perpetrator or unknown person on a specific location on that tape, as in [55], or multiple fingerprints left at a crime scene at different points in time), it can be difficult to follow this method as it offers no guidance on how to model complex conditional dependencies that come along with complex problems. This method is aimed at a particular sub-domain of the forensic domain, but it not straightforwardly

employed in other domains.

To conclude, we have examined the method from Section 4 on the advantageous BN graph properties established in Section 7. This method is aimed at BN graph construction for cases from a sub-domain of the forensic domain, namely that of forensic biology traces, which are evaluated under activity-level forensic hypotheses. The observations we have made for this method relate to this sub-domain only, but in Section 11 we deal with the forensic and legal domains more broadly.

In following this method the resulting BN graph reflects the new properties, with the exception of property P5.1. We argue that Properties P3 and P4 are not adhered to as much, because this method pays little attention to minimising the computational complexity of the model. Rather than computational feasibility, the focus in this method is the representational value of the BN graph: the resulting BN graph is seemingly more detailed than strictly necessary because of separate nodes for each activity, multiple nodes between hypotheses and evidence and even the inclusion of non-disputed activities. The authors argue that this is a possible loss of efficiency is worth it for the increase in the graph's capacity to communicate the reasoning considered in the case. Furthermore, the steps implicitly prescribe an arc direction as specified in property P7, but the method does not explicitly focus on the type of relationship between variables. A similar observation can be given for property P2: the way the variables are connected, the structure can allow for explaining-away or explaining-in, but the authors do not allude to this possibility. Lastly, the evaluation of the graph and iterative character of the construction process is not emphasised.

9 Guidelines

In the previous two sections we researched advantageous properties for BN graphs and how these relate to Taylor and colleagues’ template for BN construction. In this section we summarise our findings and propose a new set of guidelines for the manual construction of BN graphs in data-poor domains.

Firstly we conclude that depending on the rationale for building the BN, different graph properties are desired. There is a dichotomy between computational efficiency and representational value, and improvements in one area are often at the expense of the other. The motivation behind construction of the graph can vary depending on the phase of construction: in early attempts at designing the graph, the emphasis is often more on the representational value of the graph to allow for discussions on the included variables and representation of the independence relation amongst them. In later stages the graph can be simplified or adjust for the benefit of computational efficiency.

The template by Taylor and colleagues promotes the representational value of the BN and focuses little on the computational feasibility of the model. We prefer to provide guidance for both motivations for building a BN. As such, multiple BN graphs can be developed with these guidelines, varying in purpose.

The new properties, outlined in Section 7, serve as an inspiration for the formulation of the guidelines. The resulting BN graphs should adhere to the new properties as much as possible, and the steps in the guidelines are specified with this goal in mind.

Taylor and colleagues’ template is practical: it is a stepwise method that defines which variables to model, where they should be in the structure, and how to connect them. We hope to achieve a similar functionality, but have to keep in mind that these guidelines are not domain-specific. As a consequence we therefore develop these guidelines on a more abstract level by focusing on the roles of variables and type of relationship between them.

The template dictates a linear process of graph construction, whereas we argue that this is too simple. Properties P3, P4 and P6 for instance implicitly represent an evaluative check in which the graph can be altered, indicating at least some level of iteration. It is already mentioned in Section 2.3 that the BN construction process is iterative, and we find that the graph construction process is too. This should be reflected in the guidelines. Additionally, Taylor and colleague’s method pays little attention to evaluating the structure. More room should be given to evaluation of the (intermediate) graph.

We furthermore note that although the literature discussed in Part I of this thesis was helpful as inspiration for these guidelines, not all idioms were useful. The measurement, induction and reconciliation idiom (see Section 3.2.3) were not used in designing these guidelines, as these were not reflected in the BN graphs we studied.

Lastly, although causality takes a prominent place in BN modelling (see Section 3.2.1), not all relationships between variables are causal. The idioms reflect this, and so does property P7. We focus on more types of relationships

in our guidelines.

In the remainder of this section we propose our set of guidelines for the manual construction of BN graphs in data-poor domains. They are written for BN engineers tasked with constructing a BN with help of domain experts and aid them in the manual design of the network structure. We aim for these guidelines to be:

- General enough to be applicable to problems from any domain;
- Flexible enough to accommodate to specificities of problems from any domain;
- Specific enough to be practically useful;

In developing a BN graph, multiple versions are often developed in parallel. Our guidelines specify that this is a possibility. If at the end multiple eligible candidate BNs emerged from the steps, one can make an informed decision of which BN graph is the best candidate from the set of comparison measures for BN graphs, presented in Section 10.

The guidelines are listed first, after which they are discussed in more detail:

1. Identification and classification of the variables:
 - (a) Identify the relevant variables in the case;
 - (b) Establish the values;
 - (c) Categorise the variables into sets of hypothesis, evidence and mediating nodes.
2. Establish the network structure:
 - (a) If any, assess the type of relation between the variables from each category;
 - (b) Draw an arc if there is a direct influence, with the arc orientation based on the type of relation.
3. Evaluation of the structure:
 - (a) on computational feasibility;
 - (b) on representational value, i.e. as a visual aid providing an overview of the variables and reasoning considered in the case;
 - (c) on BN graph requirements, i.e. the fulfillment of the basic conditions for BN graphs (no cycles, correct representation of the independence relation of the variables according to the views of the domain expert).

Although the steps above suggest a linear process, one can alternate between the steps. The process is sequential in that one starts at steps 1a and ends at 3c, and generally advances from 1 to 3, but steps can be (and generally will be) revisited at any point in the method.

9.1 Identification and classification of the variables

The first task is to identify the variables to be included in the model, as well as their values. Although these two tasks are drafted as separate steps 1a and 1b, in practice these are intertwined. Decisions regarding variables (such as combining them) usually affect values. This task of obtaining a well-crafted variable set with values is often not straightforward and requires practice.

1a: Identification of the variables

The following can help in identifying the relevant variables: example:

1. The first step is to identify the hypothesis variable(s) (see Section 3.1). These can be established by asking what the problem or hypothesis to be tested is.
2. The next step is to gather information affecting the hypotheses: What precedes, causes, correlates with and/or results from the problem? These can be, but are not limited to events, actions, concepts, and other aspects, henceforth denoted as ‘determinants’.
3. Can other relevant determinants be identified that affect the variables established in the previous step?

Thought should be given to the name of the variable, so that it clearly portrays the element it represents. The variables identified in steps 2 and 3 can be classified as information variables (the symptom and background variables) and mediating variables from Section 3.1. Later in step 1, more specifically in Section 9.1, when the values are also established, the variables are categorised according to their role in the graph, such that it is easier to assess the relationships between the nodes.

In the previous sections we have seen that separate nodes are generally preferred: the template (Section 4) calls for separate nodes for each activity, and properties P1, P5.1 and P5.2 prescribe separate modelling as well. We therefore advise to draw separate nodes representing each identified determinant. We argue this facilitates probability estimation and generally enhances graph comprehensibility.

That being said, the computational feasibility of the model has to be kept in mind, and sometimes not every determinant has to be represented separately. Based on step 3 of Taylor and colleagues’ template (Section 4) and our research on the consequences of combining nodes (see the discussion of property P1 in Section 7), it is possible to collapse multiple variables into a single variable if:

1. The influences between them and other variables are identical or at least similar;
2. Their combined effect is not larger than each effect individually, or is not relevant;

3. Combining them does not lead to a loss of representational value of the graph, and;
4. There is no influence between one of the to-be-combined variables or its values and another variable (separate from the other variable(s) with which it will be combined).

This way, a structure including a series of nodes highly similar with respect their influences and with complex dependencies between them can be avoided. To illustrate, suppose smoking increases blood pressure, and so does alcohol. Suppose these effects are similar, in compliance with condition 1. If for the issue at hand it is not relevant how much how the combined effect of smoking and alcohol influences blood pressure or there their combination does not lead to an increased effect on blood pressure (2), they can be modelled in the same node. By naming the node ‘smoking and/or drinking alcohol’, which can be binary, no representational value is lost (3). This decision is up to the BN engineer, who should confer with the domain expert. Now suppose there is another variable which is caused by alcohol but not smoking according to the domain expert, such as ‘memory-loss’. This is not in accordance with 4, and collapsing the two variables is problematic: probabilities have to be estimated for memory-loss given smoking and/or drinking alcohol. In this situation, keeping the nodes separate is the best option.

1b: Establishing the values

When identifying the variables, the values for each variable have to be determined. First of all, the values of a variable must be mutually exclusive and exhaustive to qualify for inclusion into the model.

Care should be taken when establishing the values for a node: more values means more parameters to be estimated, i.e. a higher computational complexity. It is advisable to only include the states for a variable relevant to the current problem in the eyes of the BN engineer and domain expert [6]. Suppose for example that a variable ‘Temperature’ can take the values {high, medium, low}. If for the problem at hand the only important value of the ‘Temperature’ variable is whether it is high or not, there is no need to include ‘medium’ as a value. Note that now the value ‘low’ covers both the states ‘medium’ and ‘low’, which can be confusing. It is advisable to change the values, for example to {high, other} where ‘other’ signifies anything below ‘high’. It often suffices to keep nodes binary: the ‘Temperature’ node from the example can also be changed to ‘High temperature’ and be made binary. Here we see that value determination is closely connected with variable formulation, since establishing the values can bring along changes in the variables as well.

Ideally, there should not be any variables whose values are mutually exclusive with those of another variable. Suppose a BN engineer and domain expert come up with two binary variables, ‘Age: <50 years’ and ‘Age: ≥50 years’: their values are mutually exclusive, since if ‘Age: <50 years’ is true, ‘Age: ≥50 years’

must be false. It can be valuable to convert them into a single variable, say ‘Age’, with values $\{<50, \geq 50\}$, to reduce the complexity of the model.

In some cases this principle can be violated, as at times it is desired to keep the two variables separate for reasons of increasing the representational value of the structure and distinct chains (see the discussion on property P1 in Section 7). We advise to keep mutually exclusive variables distinct when:

- Their individual content promotes the representational value. By this we mean that separate variables can be informative visually and can communicate what variables were considered in the model. It might not be valuable to have two binary variables ‘Low temperature’ and ‘High temperature’, but it might be valuable to have variables ‘Cause of death: accidental’, ‘Cause of death: intentional’, and ‘Cause of death: natural’, as opposed to one variable ‘Cause of death’ with values $\{\text{accidental, intentional, natural}\}$.
- Other variables have an influence on only one of the variables, in other words: there is need for distinct chains. Continuing the previous example, separating variables avoids having to estimate the probability of the presence of a bloody knife given an accidental or natural cause of death.

To enforce mutual exclusivity between these nodes, a constraint node can be included if necessary (see Fenton, Neil and Lagnado [14] as discussed in Section 7).

1c: Categorising the variables

Once the variables have been identified and the values established, the resulting set of relevant variables is denoted as \mathbf{V} . The variables in \mathbf{V} should be categorised, so that $\mathbf{V} = \mathbf{H} \cup \mathbf{E} \cup \mathbf{M}$, where:

- H:** The set of variable(s) in \mathbf{V} which can be classified as hypothesis variable, i.e. those of which the probabilities are desired (see Section 3.1).
- E:** The set of evidence variables which are observed in the case to be modelled or which are expected to be observed. This set consists of the information variables established in Section 9.1, but is renamed ‘evidence’ variables to clarify that these are or can be observed.
- M:** The set of mediating variables. These are generally not observable, but provide information on hypothesis or evidence variables or how these are connected.

Classifying variables is a similar concept to the causal approach described by Kjærulff and Madsen ([26], see Section 3.2.1). We however refrain from stating a general causal relation between groups of variables, but advise to examine the relation in more detail for each pair of nodes.

In concluding step 1 of our guidelines we would like to mention that it is possible to obtain different sets of variables for the same case, for example due

to identifying variables at different levels of granularity. To illustrate, one set of variables can include nodes ‘Soil’, ‘Groundwater’ and ‘Surface water’, whereas the other summarises all these in one node ‘Environmental conditions’. This of course leads to different BN graphs, and the comparison measures in Section 10 can assist in assessing which is preferred for the task for which it is built.

9.2 Designing the graph

In this step the variables are connected by arcs to design the network structure. Step 2 of our guidelines is quite demanding of the BN engineer and domain expert: they have to take into account many construction features simultaneously. When in doubt about the arc presence or arc direction, we advise to construct multiple candidate BN graphs in parallel, and determine the best one with use of the comparison measures from the next section.

To draw arcs, the relations between nodes have to be identified. To maximise the value of the BN graph as a representational tool, we focus on the types of relations discussed in Sections 3.2.1 and 3.2.2.

Assess the relation between each pair of nodes $V_i, V_j \in \mathbf{V}$. Inspired by Taylor and colleagues’ method, we start with the hypothesis variable(s), then the evidence variable(s) and lastly the mediating variable(s).

Specifically, start with the variables $V_i \in \mathbf{H}$ and $V_j \in \mathbf{E}$: Identify whether there is a *direct* influence between V_i and V_j . If so, an arc should be drawn between them. Determine the arc direction based on the relation between the nodes:

- If V_i is a cause for V_j , draw an arc $V_i \rightarrow V_j$.
- If V_i is a sub-attribute of V_j , draw an arc $V_i \rightarrow V_j$.
- If V_i represents an event that takes place before V_j , draw an arc $V_i \rightarrow V_j$.
- If V_i is somehow associated with V_j in another way than mentioned above, the arc direction is arbitrary. Make a choice while considering:
 1. which arc direction leads to the easiest probability estimation;
 2. which arc direction leads to the lowest number of parents;

The above also holds if V_j is a cause, sub-attribute, etc. for V_i . Note that in many cases, there is seldom a *direct* influence between a hypothesis and evidence variable: oftentimes mediating variables appear in between.

Do the same for:

1. $V_i \in \mathbf{H}$ and $V_j \in \mathbf{M}$
2. $V_i \in \mathbf{M}$ and $V_j \in \mathbf{E}$
3. V_i and $V_j \in \mathbf{E}, i \neq j$
4. V_i and $V_j \in \mathbf{M}, i \neq j$

We consider it good practice to place root nodes at the top of the BN, and place the nodes in such a way that arcs are directed downward (relatively) vertically, if possible. This mirrors how Taylor and colleagues have created their template, and the idioms have a similar appearance. This way, the resulting BN will end up having a layered architecture. Uniformity in this area can allow for easier reading of the graph.

In order to increase the interpretability of the BN graph, it is possible to assign a label x to the arcs in the structure denoting the type of relationship between to variables ¹⁴. We propose a labelling of an arc $\overset{x}{\rightarrow}$, $x \in \{c, t, d, a\}$, where x signifies whether a link is causal (c), temporal (t), definitional (d), or associative (a). Other labels can be added in case other relations are inherent to a particular domain.

If the influence between variables is uncertain, there are several possibilities:

1. Assume (conditional) independence and refrain from drawing an arc: this leads to a slight loss of information, but since this information is uncertain, a simpler structure is preferred. This is especially desirable if the motivation for building the BN is computational efficiency.
2. Add a mediating node capturing the uncertainty between parent(s) and child: for example in case of disputed paternity, a putative father and an alternative father can both contribute to a child’s genotype, but not at the same time. A mediating node simulating the true father can be added in between the parent and child nodes to resolve this uncertainty. This option often benefits the representational value of the graph.

Whilst at this step, the following modelling options (inspired by those from Taylor and colleagues’ method) can be taken into account:

- There can be more than one mediating node on the same chain. Caution has to be taken: the longer the chain, the more likely it is that the impact of a node early in the chain on a latter one is weaker, and the more uncertainty is propagated through the chain [6].
- Allow for intercausal interaction in the structure: in assessing the relations between nodes, keep in mind whether there is any interaction for which the graph has to be adapted to reflect this interaction (Section 2.2.1). To illustrate, suppose two direct causes¹⁵ V_i and V_j have been identified for an effect V_k . In case these causes are in competition, i.e. each cause individually can account for the occurrence of V_k and they are mutually exclusive, they could explain each other away, which the structure should allow for. From the steps above, two arcs $V_i \overset{c}{\rightarrow} V_k$ and $V_j \overset{c}{\rightarrow} V_k$ arise, such that a converging connection emerges. Keep in mind that the converging connection need not involve direct arcs for explaining away to happen (although the effect might be weaker the longer the chain).

¹⁴This idea is inspired by Vlek and colleagues [52].

¹⁵Again, we note that a causal relationship is not strictly necessary for intercausal interactions to take place, but for the sake of the explanation we assume a causal relationship.

- The chain between hypothesis and evidence must be active, otherwise the evidence cannot influence the hypothesis. This step follows property P6. Active chains are desired unless it is certain that observing a variable cancels the influence of evidence on hypothesis. This can for example be the case when a contract is found (evidence) that is the result of a hypothesised deal between two parties. If however there is evidence that the document is a fake, the impact of the document can be undone by blocking the chain using this evidence. We strongly advise against using inactive chains in most cases, because the uncertainty regarding the evidence that blocks the chain is not incorporated.

9.3 Evaluation of the structure

Our research suggests that in constructing the graph, the following have to be taken into account: basic functionality, computational complexity, and representational value. None of the methods discussed in Part I include an explicit step to evaluate how the BN graph performs with respect to these goals. This is why our method includes an explicit step for evaluation of the structure.

The evaluation steps described below are drawn up in this order with a reason, but this order is not strict. The general idea behind this ordering is that in assessing the complexity of a model (step 1), there is a chance the model will look very different due to the removal of arcs and/or nodes, insertion of mediating nodes or reversal of arrows. Afterwards the model network should be inspected on its role as a communication tool (step 2): does the graph still provide a comprehensive visual overview of the case? These two steps can be assessed alternately until a balance is found between complexity and representation that suits the BN engineer and domain expert.

The evaluation of the structure should always conclude with analysing whether the structure fulfills the basic requirements of a BN graph (step 3): correct representation of the independence relation and acyclic. If any adjustments are made to the structure in this last step, we advise that the graph is re-evaluated, to make sure it still adheres to the intentions of the BN engineer and domain expert. In case the guidelines lead to multiple BN graphs, the candidate graphs can be evaluated on the comparison measures (step 4) proposed in the next section (Section 10).

The (initial) graph should be evaluated on the following:

1. Check the computational complexity of the structure:
 - (a) Check the number of parents for each node: strive to keep this number to a minimum. Divorcing can initially be used to reduce the number of parents, and removal of weak arcs can be done once (initial) parameters are specified.
 - (b) Critically assess the (number of) values for the nodes: is every value necessary, or can values be excluded or contained in another value?

- (c) (After (initial) parameterisation) examine the nodes: does every node that is included in the model have a significant influence on at least one other variable, or can a node be excluded (whilst adjusting the parameters for other nodes)? Sensitivity analyses (Section 2.3.1) can be used to this end.
2. Check the representational value of the structure:
 - (a) Check whether introducing mediating nodes can help clarifying the model structure with respect to its representational value, or facilitate easier probability estimation.
 3. Check the functionality of the structure:
 - (a) Verify the conditional (in-)dependencies represented by the structure: for each pair $V_i, V_j \in \mathbf{V}$, check whether their conditional (in-)dependency is correctly mirrored in the structure. The structure should reflect the independence relation of \mathbf{V} as considered by the domain expert.
 - (b) Check for cycles: cycles are not allowed in a BN graph. To ‘break’ a cycle, arcs can be reversed or removed.
 4. In case multiple candidate BN graphs emerge from the method: base the final decision on the comparison measures (discussed in Section 10).

10 Comparison measures

This thesis is aimed at providing aid in constructing the BN graph by hand. As mentioned in the guidelines, multiple candidate graphs can emerge, and therefore part of this aim is to formulate measures to allow for a comparison of BN graphs and making an informed decision.

These measures are intended for the comparison of two or more different BN graphs which model the same problem. As mentioned before, different BN graphs can all be correct in the sense of adequately representing the knowledge and views of the domain experts, but some might be more appropriate or suitable than others for the case at hand or depending on the rationale for designing the BN. If during the construction of the graph multiple structures emerge, the decision which graph is favoured over others can be based on these measures. The last step of the newly proposed guidelines (Section 9) incorporates the use of these measures.

The guiding principles from the literature (Part I) and the new properties (Section 7) again serve as inspiration in the formulation of these measures. We assume that the BN graph candidates qualify as a BN graph in that they are acyclic digraphs for the case at hand.

10.1 Development of the measures

We set out to formulate concrete comparison measures, preferably in the form of quantifiable statements. We based initial versions of our measures on the new properties, as these capture advantageous BN features.

The earliest variants of the measures were formulated in more absolute terms, which we illustrate here. In comparing the selected BN graphs, the “best” graph is the one which adheres to a number of measures, such that this graph:

1. Has the lowest number of parameters to be estimated;
2. Does not contain redundant arcs;
3. Does not contain nodes with an in-degree of four or more;
4. Does not contain more than five consecutive nodes between the hypothesis and evidence nodes;
5. Contains only binary nodes;
6. Contains separate nodes for each alternative cause for a common effect;
7. Contains separate nodes for each hypothesis;
8. Only includes arcs that represent a causal, temporal or definitional direction.

This quickly proved too radical, as these measures often lead to the premature exclusion of good BN candidates. All criteria are absolute, either *not* allowing a particular characteristic or *only* allowing a specific property. Absolute numbers such as in criteria 3 and 4 are based on observations from the database. These numbers might be concrete and easily established, but such a number cannot be straightforwardly generalised across all domains. Additionally, there is the possibility that all BN graphs in the comparison are excluded when neither of them adheres to these criteria. We approached this problem by mitigating these criteria to “contains the lowest/highest number of ...”. Other issues remained, however.

One such issue is inherited from the new properties: the measures are inherently inconsistent. Measures 5 and 6 contradict measures 1 through 3, making it difficult for a graph to be in line with all criteria in the list. We tried to solve this by coming up with priority ordering: Not every measure had to be obeyed, and in choosing the most suitable BN for the case at hand one can prioritise particular measures over others. Measures 5 and 6 for instance can be prioritised over the others in case the BN is used as a visual tool. This was an improvement, but still the measures were not satisfactory.

Lastly and most importantly, we found that the measures did not achieve what we wanted them to achieve. The motivation behind the first five measures is ensuring that the chosen BN graph provides the optimal computational complexity (whilst still adequately representing the case at hand), but this is not guaranteed with these measures. There are for example situations where a graph with a higher number of nodes and arcs (and hence, a higher number of parameters) is preferred over more compact networks (as in Section 11.4). Measures 6 through 8 were drafted with the goal of choosing the BN with the maximum representational value. The measures merely favour networks with separate nodes and interpretable arc directions, and do not prohibit that a visually overly-complex network is favoured over a more compact but clearer one.

We tried to formulate a set of detailed level, which did not work well. We concluded that we had to distance ourselves from formulating concrete measures and examined comparable aspects on a more abstract level: in what situations are particular BN graph properties preferred? The division in motivations for the BN graph returned: using the BN to perform probabilistic inference and therefore aiming for the best computational feasibility leads to different modelling choices than when using the BN graph as a visual tool to communicate the reasoning considered in the modelled case. The best computational feasibility and representational value are sums of several components. We formulated our measures in such a way that they capture these components and in this way provide a good starting point of aspects to consider when comparing BN graphs.

10.2 Final measures

The measures are as follows: For x BN graph candidates ($x \geq 2$), choose the BN graph which:

1. Is computationally most feasible, based on:
 - (a) the number of parents per node.
 - (b) the number of values per node.
2. Has the most representational value:
 - (a) separate nodes are included for each cause, hypothesis and piece of evidence.
 - (b) the arcs reflect a causal, temporal or definitional direction.

10.3 Measure 1 - Computational complexity

As mentioned before, comparing BN structures on computational complexity is not as straightforward as comparing the absolute number of parameters to be estimated¹⁶, which is dependent on the number of nodes and number of values. In some cases, a structure that is more extensive in terms of the number of nodes is preferred over a more compact graph with a higher number of nodes with many parents (especially if the nodes are non-binary valued). We therefore advise to assess the number of parents per node for each graph and let this weigh in on the decision.

In addition to the number of parents, the number of values determines the complexity of a BN structure. Generally the BN graph that contains only binary nodes or contains the highest number of binary nodes (excluding deterministic nodes) is preferred.

10.4 Measure 2 - Representation

A BN graph should demonstrate a level of transparency of the reasoning used in building the structure. The subjectivity of this characteristic is difficult to quantify into a comparison measure. We argue that representing variables in separate nodes as much as possible is a good starting point. BN graphs with this property are preferred. In addition to a separate representation, the arc orientation can help in showing the direction of reasoning, in particular for Markov equivalent graphs. Structures with (the highest number of) causal, definitional or temporal arcs increase the interpretability of the graph, and are therefore preferred over others.

10.5 Priority ordering

As measures 1 and 2 contradict each other, a priority ordering should be given to these measures, which is dependent upon the preferences BN engineer, the domain expert and the held from which the problem to be modelled originates. In some domains where complex and extensive BNs are the order of the day,

¹⁶Note that two Markov Equivalent structures for one set of variables with set values share the same set of immoralities and therefore have the same number of parameters.

computational efficiency might be higher on the list of properties of a BN graph than for other domains with simple and small models.

An ordering can be given as follows: measure 1 \leq measure 2 denotes that measure 2 is preferred over measure 1. Adding this priority ordering to the documentation of how model was built, with motivations why, could help in communicating the goals for the BN (i.e. as a communicative tool or to compute probabilities (most) efficiently).

Part III

Case Study

11 Evaluating the guidelines and comparison measures for the forensic and legal domains

In Part II we proposed a new set of guidelines for the manual construction of BN graphs with the help of domain experts, as well as a set of measures for the comparison of multiple BN graphs. As we have aimed to keep these general, it is worth examining how they relate to specific data-poor domains. In Section 8 we already provided a careful examination of Taylor and colleagues' template with respect to the new properties that form the basis of our guidelines. The method is aimed at a sub-field of the forensic domain (i.e. biological (DNA) traces under activity-level forensic hypotheses), whereas now we will look at the forensic domain in general with forensic hypotheses at any level (if any), as well as the legal domain. The legal domain is included since BNs for legal cases bear a resemblance to BNs from the forensic domain: they include crime-level forensic hypotheses and deal with (forensic) evidence.

In this section we applied the information obtained in the interview that was conducted with a forensic scientist with DNA expertise at Netherlands Forensic Institute (NFI) and examined whether the BNs from the forensic and legal domains adhere to our new guidelines and comparison measures, and if there is room for extending and/or adjusting them specifically for these domains. To this end we reassessed the database of BN graphs from the forensic and legal domains (see Section 6) on whether they are in line with what the guidelines and measures prescribe. We examined how the BN graphs are established with respect to the three main steps from the guidelines, and the measures:

1. *Variable identification:* we assessed what variables and values were included in the graph given the case description, and how these are established. We examined whether the different options for modelling variables and values are reflected in the BNs.
2. *Graph development:* we analysed the choices in arc presence, arc absence and arc direction. In addition we looked whether any of the authors made use of the specific modelling options mentioned in this step of the guidelines.
3. *Structure evaluation:* we surveyed whether any of the authors mention that they paid attention to some form of evaluation of the structure and researched whether the authors had a tendency to prefer enhancing the representational value of the graph or optimising the computational complexity.
4. *Comparison measures:* We tried comparing multiple BN structures for the same case, but could only find one instance where this was possible.

Few authors provide motivations for their choices, but conclusions can be drawn from the way the graphs are built. The BN graphs are used as examples to illustrate the concepts from the guidelines.

We discuss the BNs from the literature in relation to the three main steps from the general guidelines, i.e. variable identification, graph design, and structure evaluation.

11.1 Variable identification

A crucial difference with the guidelines and forensic BN graph conventions is the inclusion of a single node for the pair of forensic hypotheses (see Section 4.1). Our guidelines suggest separate nodes for each hypothesis, whereas in forensic BNs these are collapsed into one node to provide a direct reading of the LR. Nearly all forensic BNs from the database possess this ‘forensic hypothesis node’¹⁷ or at least a decision node. Indeed the inclusion of a single node facilitating the direct reading of the probability of interest is even explicitly advocated by Dawid and colleagues [11]. The first step in the template from Taylor and colleagues prescribes modelling the forensic hypotheses in human DNA cases as distinct values of a single node. In the discussion of this first step of the template in Section 8 we argued that this hypothesis node is beneficial for this specific sub-domain, but it appears that this can be generalised to the domain of forensic science. The proposed guidelines could be adapted so as to include a single node for the forensic hypotheses for BNs built in the domain of forensic science: steps 1a and 1b of the guidelines could contain an additional instruction to draw a single node for both forensic hypotheses.

Some authors make use of the grouping of variables and/or values. In their BN shown in Figure 25 in Appendix A, Taylor and colleagues [47] group together two pieces of evidence, namely two tapelifts of the inside and outside of an item clothing, and regard these as one large tapelift. Both tapelifts revealed quantities of DNA of the suspect, but only on one tapelift low levels of DNA from another person was detected. Despite the fact that the DNA from this other person was only found on one tapelift, the authors argue that there is little evaluative difference between regarding the presence and absence of DNA in these bordering surfaces and considering the presence of DNA for the general area (i.e. the whole area of the item of clothing from which the tapelift evidence is gathered).

Although the case modelled by Szkuta and colleagues [44] (see Figure 24 in Appendix A) involves two victims, the authors have grouped both these victims when considering the presence and transfer of their DNA on pieces of evidence. For the questions in the present case (i.e. whether a suspect ‘F’ was present), and who held a piece of evidence (‘F’, another suspect ‘S’, one of the victims or an unknown individual) the impact of the victims individually is not relevant.

¹⁷Exceptions being Fenton et al [18] (although they do include a constraint node to enforce mutual exclusivity between the forensic hypotheses (see Section 7)), Oosterman et al [36], Smit et al [43], Vlek et al [51], and De Zoete et al [59].

Note that it would have been possible for this BN to construe a single node ‘DNA on didgeridoo’ with mutually exclusive¹⁸ and exhaustive values {F, S, V, U}. This is however not desired for the sake of distinct chains, as well as representation. To illustrate, the variable ‘Background DNA’ for example considers the presence of DNA of an unknown prior to any offence and is only relevant to ‘U’. Collapsing the nodes would force the probability estimation of each individual’s DNA on the piece of evidence given the background presence or absence of an unknown person’s DNA. In addition, the persons considered in the case are directly observable from the structure of the graph. Szkuta and colleagues model mutually exclusive causes as separate nodes, but do not use a constraint node (see Figure 21) to enforce this mutual exclusivity. From the interview [27] it followed that explicitly enforcing mutual exclusivity is not typically performed. It turns out that only Fenton and colleagues [18] and Vlek and colleagues [51] make use of this constraint node in their BN, see Figures 36 and 38 in Appendix A.

Szkuta and colleagues furthermore consider the transfer, persistence and recovery of DNA together in the binary nodes ‘Transfer/Persistence [...]’. The BNs by Taylor and colleagues [47] and De Ronde and colleagues [42] (Figures 25 and 26 in Appendix A), both created using Taylor and colleagues’ template (see Section 4), also consider these together. Wieten and colleagues [55] (Figure 27 in Appendix A) favour to include these separately. This last approach is encouraged by the guidelines, as it allows for easier probability estimations and increases the comprehensibility of the BN graph. We note however that this is a domain-specific issue for which we do not possess the knowledge to determine which way is best: it might well be possible for the domain expert to consider these variables together.

11.2 Network structure

The majority of the papers on forensic and legal BNs we examined have no documentation on the realisation of the BN structure. This complicates relating the structure to our guidelines, as we cannot be certain of the considerations behind the modelling choices.

Source-level BNs tend to include fewer ‘causal arcs’, and more definitional arcs. Source-level BNs concerning the identification of donors of DNA, such as the paternity network by Dawid and colleagues [11] or the mixed DNA trace network by Mortera and colleagues [34] (Figures 28 and 32 in Appendix A), include separate nodes for the paternal and maternal genes that *define* a genotype. The source-level BN by Taylor and colleagues [46] includes an example of a temporal arcs: the arc directed from the node ‘HemaStix results’ to ‘HemaTrace results’ reflects a temporal relation, since the first is carried out before the other. De Wolff and colleagues [57] initially drew an arc between two tests included in their case in the same way, which was later removed (discussed further in Section 11.3).

¹⁸The authors make the assumption that only one person handled the didgeridoo at the time of the offence, making these values mutually exclusive.

Causal relations can be found between variables from activity-level BNs, as activities, transfer, persistence, and contamination can generally be viewed as causes for the findings in a case. Examples can be found in the BNs resulting from Taylor and colleagues template, such as one by De Ronde and colleagues [42] in Figure 26 in Appendix A. The ‘4. Fingermarks S through climbing’ and ‘5. Fingermarks S through leaning’ are causes of the red nodes that represent observed pieces of evidence. Contamination of a sample can be considered a cause of a particular outcome in analysing that sample. All BNs taking contamination into account, such as in Biedermann, Bozza and Taroni’s gunshot residue BN [4] and Wieten and colleagues tape BN [55] (Figures 31 and 27 in Appendix A), direct arcs away from the contamination nodes, directly or indirectly towards findings nodes.

There are few mentions of explicit intercausal interactions in the BNs from our database. Biedermann and colleagues parameterised their “cause-of-fire” BN such that the presence of terpenes, an organic combustible compound, within the floor explains away the presence of terpenes at the sampling point of a fire (see nodes ‘L’, ‘X’ and ‘T’ in Figure 23 in Appendix A). Even though explaining away is desired (and possibly used) in other BNs, such as the one by De Ronde and colleagues, the use of this effect is not discussed in the article. No other types of intercausal interactions could be identified in the BNs from the database.

From the interview [27] we gathered that intercausal interactions such as explaining away were not taken into account in BN (graph) construction at that particular department at the NFI. A reason for this is that the area of expertise of the interviewee is (human) DNA analysis, where explaining away is not always desirable. DNA can accumulate: a higher probability of DNA ending up somewhere due to an innocent activity then does not automatically mean that this person did not (also) perform a criminal activity resulting in the DNA being there. The simultaneous occurrence of both these activities is possible, and in both instances DNA could have been left. One activity explaining the other away is then not desired.

For this sub-domain of the forensic domain, the suggestion of considering intercausal interactions, or rather, explaining away specifically, from our guidelines might not be relevant. Other sub-domains however, such as fingermarks, might benefit.

The BNs from the database contain some examples of incorporating more mediating nodes to enhance the interpretability of the graph. Biedermann and colleagues [2] for example, stress that nodes ‘A: presence of gasoline after the fire’ and ‘S: presence of gasoline in the sample at the time of analysis’ can be eliminated from the graph (see Figure 22 in Appendix A, such that nodes ‘G: presence of gasoline before the start of the fire’ and ‘E: gasoline detected’ exhibit a direct dependency. For the sake of transparency, nodes ‘S’ and ‘A’ are included in this BN. They are excluded in subsequent networks from the same authors [3] built from this network (see for example Figure 23 in Appendix A), because the uncertainties captured by these nodes (namely regarding the environmental conditions in the retention of the combustible liquid and contamination of the

sample) were not relevant to these cases.

Mediating nodes are used in some models to enhance the computational feasibility and representation. Fenton and colleagues [18] for example, introduce a mediating node to summarise the capability and motive of the defendant, shown Figures 34 and 35 in Appendix A. De Wolff and colleagues [57] summarise the presence of sources other than saliva of a particular enzyme into a mediating node (see Figure 30 in Appendix A).

11.3 Evaluation of the structure

In this section we discuss whether the authors of the articles from which the BNs originate devote attention to any of the areas the evaluation step of our guidelines considers. We notice that the BNs from the database are quite simplistic. To illustrate, the smallest network contains just three nodes, the largest 57, but most contain between 10 and 14 nodes. The number of head-to-head nodes is not very high, nor is the number of parents for a single node. The focus on computational complexity is not high, and few authors mention trying to optimise their BN graph in terms of feasibility.

Most BNs tend to include more nodes to increase the transparent communication of the variables considered and reasoning involved. We already mentioned Biedermann and colleagues' choice for including nodes 'S' and 'A' in Figure 22, which are not strictly necessary but help in understanding the reasoning done in the case. Taylor and colleagues' BN [47] from Figure 25 also includes variables which are always in a particular state and do not influence the hypothesis variable. This follows from their method, which prescribes to include nodes for non-disputed activities. Some networks contains nodes with a large number of parents, e.g. that of Langford and colleagues [31], Figure 37 and of De Wolff and colleagues [57], Figure 30. Nothing about the complexity of inference is mentioned in these articles however.

The tendency to include more nodes was echoed by the forensic practitioner in our interview [27]. Motivations included clearly communicating the considerations in developing the model and easing the task of probability estimation. Their focus remains on representation, and computational feasibility does not play a role in forensic biology cases (within this department at the NFI at least). It seems safe to conclude that at the moment, in forensic and legal BN modelling, representation has priority over computational feasibility.

Regarding the independences captured by the graph, some authors motivate their modelling choices. Aitken, Taroni and Garbolino [1] (Figure 39 in Appendix A) for example discuss the presence of the arc between nodes 5 and 6, i.e. the transfer of biological material from victim to suspect in node 5 and vice versa in node 6. If it is known whether contact was established between the victim and the suspect (node 11), nodes 5 and 6 do not become independent because information on the success or failure of transfer from victim to suspect is still directly relevant for the probability of the other transfer. De Wolff and colleagues [57] mention an initial dependency between two tests (nodes 7 and 8 in Figure 30): a positive result for the first test could predict a positive result

for the second test, or the sample could be wasted during the first test, influencing the outcome of the second test. Either way, a dependency was assumed and an arc was drawn, and a node indicating whether or not the first test was performed was added. In the subsequent case study however it appeared the presumed dependency was not significant, and the arc and additional node were removed. It should be noted that although this removed the direct dependency between the two tests, an active chain remained (nodes 7, 6 and 8) and an indirect dependency might still exist.

11.4 Comparison measures

Most authors of the articles in our database neglected to include any primary attempts or other candidate BNs/BN graphs, but we established that representation generally has priority over computational feasibility within the forensic and legal domains. We base this on the fact that some BNs from the database include a large number of parents, and no divorcing was performed to minimise that number (e.g. 8 parents in the BNs from Langford and colleagues 37 and De Wolff and colleagues 30, 6 in Vlek and colleagues 38). Additionally, the interview [27] revealed that computational complexity was not an issue for the BNs constructed in this department, but rather that BN engineers concentrated on conveying the reasoning involved in the cases by including more nodes (even if their inclusion can be avoided).

Mortera and colleagues [34] did document the development of their BN graph by demonstrating two attempts for the same case in their article. They argue that the first BN graph, shown in Figure 32, is visually intuitive, but is computationally inefficient. The network contains 18 nodes, but also includes a node with at least 8 values¹⁹ and 7 parents (of which most are multi-valued as well). The size of the CPT for this node is large, 124416 entries at the very least. The authors therefore built a different network, displayed in Figure 33, which is the preferred network. While containing more nodes (46), performing calculations with this BN is computationally more feasible: the large table has been reduced to 768 entries. This network still includes separate nodes for each determinant, yet appears more complex graphically.

In conclusion, unlike most other authors we examined, Mortera and colleagues based their decision on the computational feasibility: they assessed the number of parents and parameters for each network and picked the most efficient one. Even though the first network is favoured in terms of visual comprehensibility, the second network prevailed. The comparison measures capture the criteria on which this decision was based. Suppose Mortera and colleagues would have used our measures, given that they aimed for efficient computations with their BN, measure 1 would have priority over measure 2 (i.e. $\text{measure 2} \leq \text{measure 1}$). We argue that by employing the measures they would have preferred the same BN.

¹⁹Depending on the case, different DNA markers with different numbers of alleles can be considered. For markers with a high number of alleles, the number of values for the nodes in this network can be even higher.

12 Discussion

In this section we discuss our proposed guidelines and comparison measures (Part II) and the case study regarding the forensic and legal domain BNs (Part III).

In the first place, this thesis deals with an inherently subjective topic. What constitutes a good BN graph is dependent on the personal preferences and motivations of the BN engineer and the domain expert. This brings along difficulties in designing guidelines and measures claiming to result in good BN graphs. We had to find a balance between providing instructions and offering room for the preferences of the BN engineer and domain expert whilst keeping the new properties in mind. For example, we recommended to draw separate nodes as much as possible, in line with properties P1, P5.1 and P5.2, but allowed for the option of combining variables if this was desired by the BN engineer.

Secondly, the in-depth survey of observations from the field is one-sided: only BN graphs from the forensic and legal domains were examined, and the interviewee was a forensic scientist. The risk here is that the conclusions based on these observations are partial towards these domains.

Another aspect that made the development of guidelines and measures difficult was the fact that at times the many exceptions made it difficult to develop a rule, especially in the first step of the guidelines (variable establishment) and in the measures. Many situations where the rules do not apply are context-dependent, and it proved difficult to take these into account when developing general guidelines and measures. Moreover, this complicated maintaining a clear direction throughout the guidelines, as we often had to sidestep to mention situations where the rules do not apply.

Much work in the guidelines and especially measures rests on the assumption that those involved in constructing a BN graph always incline more towards reducing computational complexity or increasing the value of the graph as a representational tool. In practice one might simply aim to achieve the optimal model without leaning more towards one motivation. We lacked practical knowledge of the motivations for building BNs and the modelling choices that go along with that, which we tried to solve by conducting an interview with a BN engineer. This was useful, but of course only shed light on BN construction in that specific area.

We aimed for the guidelines to be “specific enough to be practically useful” and have tried to achieve a level of practical applicability similar to that of Taylor and colleagues. Due to the fact that we also had to keep the guidelines general, it is uncertain whether this practicality is actually achieved. Our guidelines offer a number of good considerations for a BN engineer, but at times do not adhere to the level of stepwise graph construction as in Taylor and colleagues’ template. Because of this loss in practicality a lot of work still lies with the BN engineer and domain expert. Sometimes the guidelines are quite demanding, step 2 in particular asks to draft the network structure whilst keeping a number of aspects, which can be quite complex, in mind. On the other hand, a list of

the concrete aspects to consider is preferred over no such list at all, and in this way the guidelines do offer assistance.

The measures in particular were difficult to formulate. Earlier versions included measures that quantified a specific characteristic of the BN graph heavily based on the new properties, such as ‘choose the BN with the lowest number of parameters’. We came to the conclusion that these measures would not hold up, and eventually had to settle for less concrete measures that still provide help in making a subjective decision.

Lastly, the relation of the guidelines to the BNs from the forensic and legal domains was fruitful. It showed that most of what the guidelines postulate is reflected in the BN graphs from these domains. Sometimes the guidelines were more extensive: they include for example a step regarding the consideration of intercausal interactions, a feature that was hardly touched upon in the BNs that we studied. The case study however involved speculation, as we could not be sure of the modelling choices when these were not motivated in the articles.

Additionally we had hoped to uncover parallels between BN graphs and forensic hypotheses, but this was not the case. In other words, if characteristics specific to BN graphs with activity-level forensic hypotheses were discovered, the guidelines could have been specified further. Unfortunately we could not detect such similarities.

13 Conclusion

In current literature few forms of guidance exist on handcrafting the BN graph for a problem with help of a domain expert. The few guiding approaches that exist are either generally applicable but not very hands-on, or detailed but highly domain-specific. Practical guidelines that are not domain-specific are lacking, and their existence would be useful to increase the construction and use of BNs to tackle problems from any domain. We addressed this gap by making it the purpose of this thesis: we aimed to provide guidance in construction of a BN graph by hand with help of a domain expert.

To do so, we first gathered information by reviewing the guidance from the current literature in Part I. In Part II we used the obtained knowledge from the literature to examine what characteristics BN graphs should possess ideally (Section 7). In Section 8 we related these properties to a method for BN graph construction from Part I. The goal here was to investigate how the BN graphs resulting from this method performed with respect to the newly formed characteristics, as the method seemed a good source of inspiration for our guidelines thanks to its practical approach. In Section 9 we proposed a new set of guidelines for the manual BN graph construction for cases from data-poor domains. Furthermore we devised a set of comparison measures to compare BN structures for the same case on in Section 10. Finally in Part III we evaluated the guidelines and measures by relating them to a case study. The ‘case’ here was both the forensic and legal domains, whose BNs were assessed on whether or not they roughly followed the steps from our guidelines. We found that they generally do, but these BNs do not draw separate nodes for each forensic hypothesis, but rather incorporate a single node for the pair of them.

In concluding this thesis we revisit our research questions.

1. What are desirable characteristics for manually constructed BN graphs for problems from data-poor domains and why are these desirable?

Our efforts in extending the applicability of Timmer’s properties to BN graphs from any domain resulted in a new list of advantageous properties (Section 7.1). These properties are desirable because they either optimize the computational efficiency of the BN graph, or increase the potential of the BN graph as a visual tool. The motivations for using the BN determine when adhering to a property is beneficial: when the BN is merely as a tool for calculating the probability of interest in the most efficient way, other properties are desired than when the BN graph also functions as a communicative tool. All in all, the new properties helped in determining what constitutes a good BN graph.

2. How can these insights be combined to formulate:
 - (a) a set of practical guidelines for the manual BN graph construction with help of domain experts for a problem from any domain, and
 - (b) a set of comparison measures for analysing different BN graphs for the same problem and establishing which graph is the most suitable in a specific situation?

In Section 9 we proposed our set of guidelines. In developing them we incorporated the general approaches discussed in the literature study. Additionally, the newly drafted properties served as a goal towards which our guidelines should work to: the steps in our guidelines were developed in such a way that the BNs emerging from them would incorporate these properties. Our guidelines for example call for including explicit and separate nodes for hypotheses, evidence and causes, and promote arc direction as stated in property P7.

Furthermore, we aspired to ensure a practical applicability by designing the steps with the practical method from Taylor and colleagues in mind. As we examined how the BNs from this method scored on the new properties, we knew how to implement elements from this template and what steps to avoid. To illustrate, in assessing the relationships we start with those between the hypothesis node(s) and evidence nodes, mirroring the template. On the other hand, the template does not explicitly allude to the possibility of explaining away or any type of intercausal interaction. Because we feel this is a powerful characteristic of BNs, we incorporated it in our guidelines.

We proposed our comparison measures in Section 10. We argued that these provide some assistance in the determining which BN is preferred in particular situations, but not as much as we had hoped. It proved challenging to draft a more concrete set of measures due to the difficulties in quantifying subjective intricacies of BN graphs. To give an example, separate nodes for example generally improve the visual interpretability of a network structure, but demote it if there are too many. We formulated the comparison measures on a more abstract level and provided guidance in choosing the most suitable graph for two specific motivations for modelling the BN graph.

3. How do the guidelines and comparison measures from research question 2 relate to actual manually constructed BN graphs from the forensic and legal domains?

We evaluated the guidelines by examining whether BN graphs from the forensic and legal domains displayed the directions from the guidelines. The BN graphs are largely in line with what the guidelines prescribe, with the main exceptions of the inclusion of a single node for the forensic hypotheses and little focus on intercausal interactions. The measures were more difficult to evaluate, as the literature provided few examples of multiple BN structures for the same case. We argue that based on the overall level of detail of the BN graphs, computational complexity is often not yet an issue for these domains and the focus is on the representational value of the graph.

For future research, we suggest it would be valuable to apply the guidelines and comparison measures to an actual case. When putting them to use, any gaps or inconsistencies can be identified and adjusted. It could also lead to fine-tuning of the steps: more concrete modelling techniques, such as (leaky) noisy-OR and OOBN modelling, could be incorporated. Further proto-typical examples explaining how to handle particular situations would also be useful. Lastly, we suggest the performing a similar assessment of the guidelines for other

data-poor domains as a case study, such as the one we did for the forensic and legal domains. The guidelines could be adapted or extended to provide the best instructions for BN graph design for these specific domains.

References

- [1] Aitken, C., Taroni, F., & Garbolino, P. (2003). A graphical model for the evaluation of cross-transfer evidence in DNA profiles. *Theoretical Population Biology*, 63(3), 179–190.
- [2] Biedermann, A., Taroni, F., Delemont, O., Semadeni, C., & Davison, A. (2005). The evaluation of evidence in the forensic investigation of fire incidents (Part I): An approach using Bayesian networks. *Forensic Science International*, 147(1), 49–57.
- [3] Biedermann, A., Taroni, F., Delemont, O., Semadeni, C., & Davison, A. C. (2005). The evaluation of evidence in the forensic investigation of fire incidents. Part II. Practical examples of the use of Bayesian networks. *Forensic Science International*, 147(1), 59–69.
- [4] Biedermann, A., Bozza, S., & Taroni, F. (2009). Probabilistic evidential assessment of gunshot residue particle evidence (Part I): Likelihood ratio calculation and case pre-assessment using Bayesian Networks. *Forensic Science International*, 191(1-3), 24–35.
- [5] Boneh, T., Nicholson, A. E., & Sonenberg, E. A. (2006). Matilda: A visual tool for modeling with Bayesian networks. *International Journal of Intelligent Systems*, 21(11), 1127–1150.
- [6] Cain, J. (2001). Planning improvements in natural resources management. In: *Guidelines for Using Bayesian Networks to Support the Planning and Management of Development Programmes in the Water Sector and Beyond*. Centre for Ecology and Hydrology.
- [7] Chen, S. H., & Pollino, C. A. (2012). Good practice in Bayesian network modelling. *Environmental Modelling & Software*, 37, 134 – 145.
- [8] Cheng, J., & Greiner, R. (1999). Comparing Bayesian network classifiers. In K. B. Laskey, H. Prade (Eds.), *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*, (pp. 101–108). Morgan Kaufmann Publishers Inc.
- [9] Chickering, D. M. (2002). Learning equivalence classes of Bayesian-network structures. *Journal of Machine Learning Research*, 2, 445–498.
- [10] Cook, R., Evett, I. W., Jackson, G., Jones, P. J., & Lambert, J. A. (1998). A hierarchy of propositions: Deciding which level to address in casework. *Science & Justice*, 38, 231–239.
- [11] Dawid, A. P., Mortera, J. Pascali, V. L., & Van Boxel, D. (2002). Probabilistic expert systems for forensic inference from genetic markers. *Scandinavian Journal of Statistics*, 29(4), 577–595.

- [12] Druzdzel, M. J., & Henrion, M. (1993). Intercausal reasoning with uninstantiated ancestor nodes. In *Uncertainty in Artificial Intelligence* (pp. 317-325). Morgan Kaufmann.
- [13] Druzdzel, M. J., & Van der Gaag, L. C. (2000). Building probabilistic networks: Where do the numbers come from? – A guide to the literature. *IEEE Transactions on knowledge and data engineering*, 12(4), 481–486.
- [14] Fenton, N., Neil, M., & D. Lagnado, D. A. (2011). Modelling mutually exclusive causes in Bayesian networks. Submitted to IEEE Transactions on Knowledge and Data Engineering.
- [15] Fenton, N., & Neil, M. (2012). *Risk Assessment and Decision Analysis With Bayesian Networks*. Chapman & Hall.
- [16] Fenton, N., Neil, M., & D. Lagnado, D. A. (2012). A general structure for legal arguments about evidence using Bayesian networks. *Cognitive Science*, 37, 61–102.
- [17] Lagnado, D. A., Fenton, N., & Neil, M. (2013). Legal idioms: A framework for evidential reasoning. *Argument & Computation*, 4(1), 46–63.
- [18] Fenton, N., Neil, M., Yet, B., & Lagnado, D. (2019). Analyzing the Simonshaven case using Bayesian networks. *Topics in Cognitive Science*, 1–23.
- [19] Flesch, I., & Lucas, P. J. (2007). Markov equivalence in Bayesian networks. In *Advances in Probabilistic Graphical Models*, (pp. 3–38). Springer, Berlin, Heidelberg.
- [20] Van Gosliga, S. P., & Van de Voorde, I. (2008). Hypothesis Management Framework: A flexible design pattern for belief networks in decision support systems. In S. Renooij, H. J. M. Tabachneck-Schijf, S. M. Mahoney (Eds.), *Proceedings of the Sixth UAI Bayesian Modelling Applications Workshop*, Helsinki, Finland.
- [21] Van der Gaag, L., & Renooij, S. (2019). *Probabilistic Reasoning With Bayesian Networks* [Course Syllabus]. Retrieved from <http://www.cs.uu.nl/docs/vakken/prob/Docs/Syllabus-INFOPROB19.pdf>
- [22] Haraksim, R., Meuwly, D., Doekhie, G., Vergeer, P., & Sjerps, M. (2013). Assignment of the evidential value of a fingerprint general pattern using a Bayesian network. In A. Brömme, C. Busch (Eds.), *BIOSIG 2013: Proceedings of the 12th International Conference of the Biometrics Special Interest Group* (pp. 99–103). Gesellschaft für Informatik.
- [23] Hepler, A. B., Dawid, A. P., & Leucari, V. (2007). Object-oriented graphical representations of complex patterns of evidence. *Law, Probability & Risk*, 6(1-4), 275–293.

- [24] Jensen, F. V., & Nielsen, T. D. (2007). *Bayesian Networks and Decision Graphs*. Springer Science & Business Media.
- [25] Kjærulff, U. B. (1994). Reduction of computational complexity in Bayesian networks through removal of weak dependences. In *Proceedings of the 10th Conference on Uncertainty in Artificial Intelligence*, pp. 374–382.
- [26] Kjærulff, U. B., & Madsen, A. L. (2008). *Bayesian Networks and Influence Diagrams: A Guide to Construction and Analysis* (2nd ed.). Springer-Verlag New York.
- [27] Kokshoorn, B. (2020, June 19). Skype interview.
- [28] Kokshoorn, B., Blankers, B. J., De Zoete, J., & Berger, C. E. (2017). Activity level DNA evidence evaluation: On propositions addressing the actor or the activity. *Forensic Science International*, 278, 115-124.
- [29] Koller, D., & Pfeffer, A. (1997). Object-oriented Bayesian networks. In D. Geiger, P. P. Shenoy (Eds.), *Proceedings of the Thirteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-97)* (pp. 334–341).
- [30] Korb, K. B., & Nicholson, A. E. (2010). *Bayesian Artificial Intelligence* (2nd ed.). Chapman & Hall.
- [31] Langford, A. M., Bolton, J. R., Carlin, M. G., & Palmer, R. (2015). Post-mortem toxicology: A pilot study to evaluate the use of a Bayesian network to assess the likelihood of fatality. *Journal of Forensic and Legal Medicine*, 33, 82–90.
- [32] Laskey, K. B., & Mahoney, S. M. (1997). Network fragments: Representing knowledge for constructing probabilistic models. In D. Geiger, P. P. Shenoy (Eds.), *Proceedings of the Thirteenth Annual Conference on Uncertainty in Artificial Intelligence (UAI-97)* (pp. 334–341).
- [33] Lauritzen, S. L., & Spiegelhalter, D. J. (1988). Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society: Series B*, 50(2), 157–224.
- [34] Mortera, J., Dawid, A. P., & Lauritzen, S. L. (2003). Probabilistic expert systems for DNA mixture profiling. *Theoretical Population Biology*, 63(3), 191–205.
- [35] Neil, M., Fenton, N., & Nielson, L. (2000). Building large-scale Bayesian networks. *The Knowledge Engineering Review*, 15(3), 257–284.
- [36] Oosterman, W. T., Kokshoorn, B., Maaskant-van Wijk, P. A., & De Zoete, J. (2015). From sub-source to source: Interpreting results of biological trace investigations using probabilistic models. *Forensic Science International: Genetics Supplement Series*, 5, e590-e591.

- [37] Panofsky, H. A., & Brier, G. W. (1968). *Some Applications of Statistics to Meteorology*. The Pennsylvania State University, University Park.
- [38] Pearl, J. (1988). *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Publishers Inc.
- [39] Pitchforth, J., & Mengersen, K. (2013). A proposed validation framework for expert elicited Bayesian networks. *Expert Systems with Applications*, 40(1), 162–167.
- [40] Renooij, S. (2001). *Qualitative Approaches to Quantifying Probabilistic Networks*. Doctoral dissertation, Utrecht University.
- [41] Renooij, S. (2010). Bayesian network sensitivity to arc-removal. In *Proceedings of the Fifth European Workshop on Probabilistic Graphical Models (PGM-2010)*, 233–240.
- [42] De Ronde, A., Kokshoorn, B., De Poot, C. J., & De Puit, M. (2019). The evaluation of fingerprints given activity level propositions. *Forensic Science International*, 302, [109904].
- [43] Smit, N. M., Lagnado, D. A., Morgan, R. M., & Fenton, N. E. (2016). Using Bayesian networks to guide the assessment of new evidence in an appeal case. *Crime science*, 5(1), 1–12.
- [44] Szkuta, B., Ballantyne, K. N., Kokshoorn, B., & Van Oorschot, R. A. (2018). Transfer and persistence of non-self DNA on hands over time: Using empirical data to evaluate DNA evidence given activity level propositions. *Forensic Science International: Genetics*, 33, 84–97.
- [45] Taroni, F., Aitken, C. G., Garbolino, P., & Biedermann, A. (2006). *Bayesian Networks and Probabilistic Inference in Forensic Science*. Chichester: Wiley.
- [46] Taylor, D., Abarno, D., Hicks, T., & Champod, C. (2016). Evaluating forensic biology results given source level propositions. *Forensic Science International: Genetics*, 21, 54–67.
- [47] Taylor, D., Biedermann, A., Hicks, T., & Champod, C. (2018). A template for constructing Bayesian networks in forensic biology cases when considering activity level propositions. *Forensic Science International: Genetics*, 33, 136–146.
- [48] Taylor, D., Kokshoorn, B., & Biedermann, A. (2018). Evaluation of forensic genetics findings given activity level propositions: A review. *Forensic Science International: Genetics*, 36, 34–49.
- [49] Timmer, S. (2017). *Designing and Understanding Forensic Bayesian Networks Using Argumentation*. Doctoral dissertation, Utrecht University.

- [50] Vicard, P., Dawid, A. P., Mortera, J., & Lauritzen, S. L. (2008). Estimating mutation rates from paternity casework. *Forensic Science International: Genetics*, 2(1), 9–18.
- [51] Vlek, C., Prakken, H., Renooij, S., Verheij, B. (2013). Modeling crime scenarios in a Bayesian network. In B. Verheij (Eds.), *Proceedings of the 14th International Conference on Artificial Intelligence and Law*, (pp. 150–159).
- [52] Vlek, C., Prakken, H., Renooij, S., Verheij, B. (2016). A method for explaining Bayesian networks for legal evidence with scenarios. *Artificial Intelligence and Law*, 24(3), 285–324.
- [53] Walton, D., Reed, C., & Macagno, F. (2008). *Argumentation Schemes*. Cambridge University Press.
- [54] Wellman, M. P., & Henrion, M. (1993). Explaining ‘explaining away’. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 15(3), 287–292.
- [55] Wieten, R., De Zoete, J., Blankers, B., & Kokshoorn, B. (2015). The interpretation of traces found on adhesive tapes. *Law, Probability and Risk*, 14(4), 305–322.
- [56] Wieten, R., Bex, F., Prakken, H., & Renooij, S. (2019). Supporting discussions about forensic Bayesian networks using argumentation. In *Seventeenth International Conference on Artificial Intelligence and Law (ICAAIL ‘19)*, June 17–21, 2019, Montreal, CQ, Canada. ACM, New York, NY, 143–152.
- [57] De Wolff, T. R., Kal, A. J., Berger, C. E. H., & Kokshoorn, B. (2015). A probabilistic approach to body fluid typing interpretation: An exploratory study on forensic saliva testing. *Law, Probability and Risk*, 14(4), 323–339.
- [58] Woudenberg, S. P., Van Der Gaag, L. C., & Rademaker, C. M. (2015). An intercausal cancellation model for Bayesian-network engineering. *International Journal of Approximate Reasoning*, 63, 32–47.
- [59] De Zoete, J., Oosterman, W., Kokshoorn, B., & Sjerps, M. (2016). Cell type determination and association with the DNA donor. *Forensic Science International: Genetics*, 25, 97–111.

Appendices

A Figures of BNs

This appendix includes several figures from the database of BNs from forensic and legal domains (the complete list of articles included in the database can be found in 4 in Appendix B). We limit the description of the BNs to only that information about the BN to illustrate ideas and substantiate the claims we make in this thesis. For a full description we refer the reader to the article from which the BN originates.




Figure 22: BN for detected residual flammable liquid (taken from [2, p. 53])

The BN in Figure 22 is used to compute the probability of different causes of fire incidents given the presence of flammable liquid in a trace. The nodes signify the following (all are binary unless stated otherwise):

- H: Cause of fire, with values {natural, technical, human action}
- M: Sampling point lies within area of original fire
- D: Presence of spilled gasoline at sampling point
- B: Background presence of gasoline at sampling point
- G: Presence of gasoline before start of fire
- A: Presence of gasoline after fire
- S: Presence of gasoline in sample at time of analysis
- E: Gasoline detected

Nodes ‘A’ and ‘S’, the authors state, can be left out to reduce the number of nodes (hence parameters to estimate), but enhances the transparency of the BN graph.




Figure 23: BN for detected terpenes, taken from [3, p. 60]

The BN in Figure 23 is similar to that in Figure 22, but models a more detailed case: the probability of a particular cause of a fire can be computed given the presence of terpenes in a sample taken from the wooden floor (matrix) identified as area of fire origin. The nodes signify the following (all are binary unless stated otherwise):

- D: Spill of combustible liquid that contains terpenes
- B: Background presence of terpenes at sampling point
- L: Presence of terpenes that do not come from the matrix
- X: The matrix contains terpenes
- T: Presence of terpenes at sampling point
- E: Detection of terpenes in fire debris sample

Explaining away can occur between nodes ‘L’ and ‘X’.

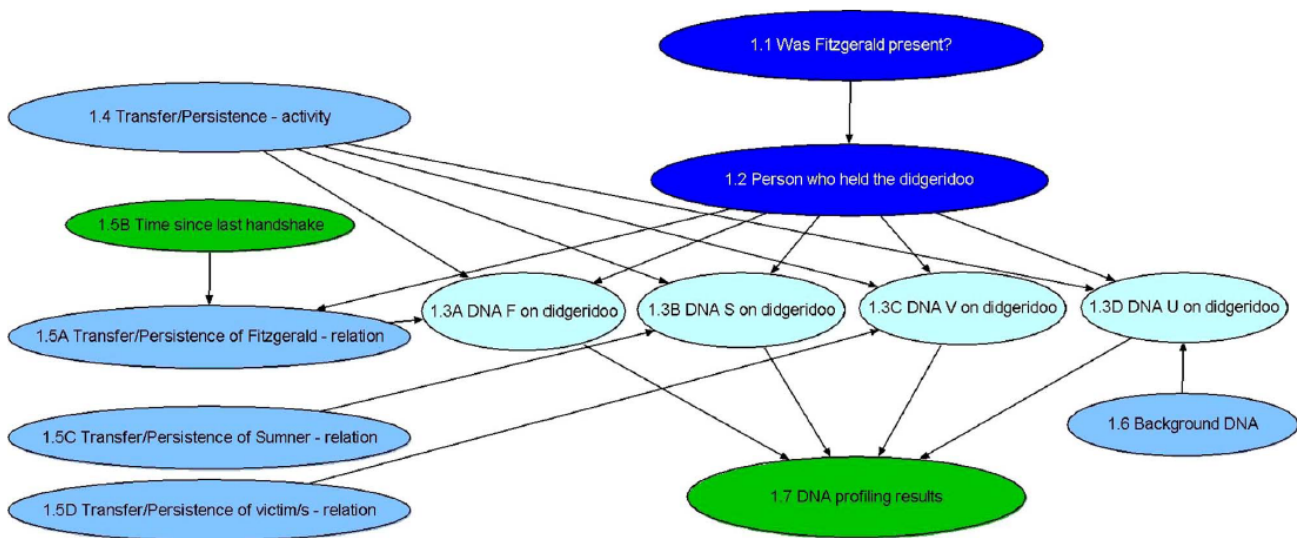


Figure 24: BN for the transfer and persistence of non-self DNA, taken from [44, p. 86]

The BN in Figure 24 models the breaking and entering into the home of two victims ‘V’, involving suspects ‘F’ and ‘S’. The two victims are grouped together under ‘V’.

The activity-level forensic hypotheses are:

H_p F was present in the house during the attack and held the didgeridoo.

H_d F was not present in the house during the attack and someone else held the didgeridoo.

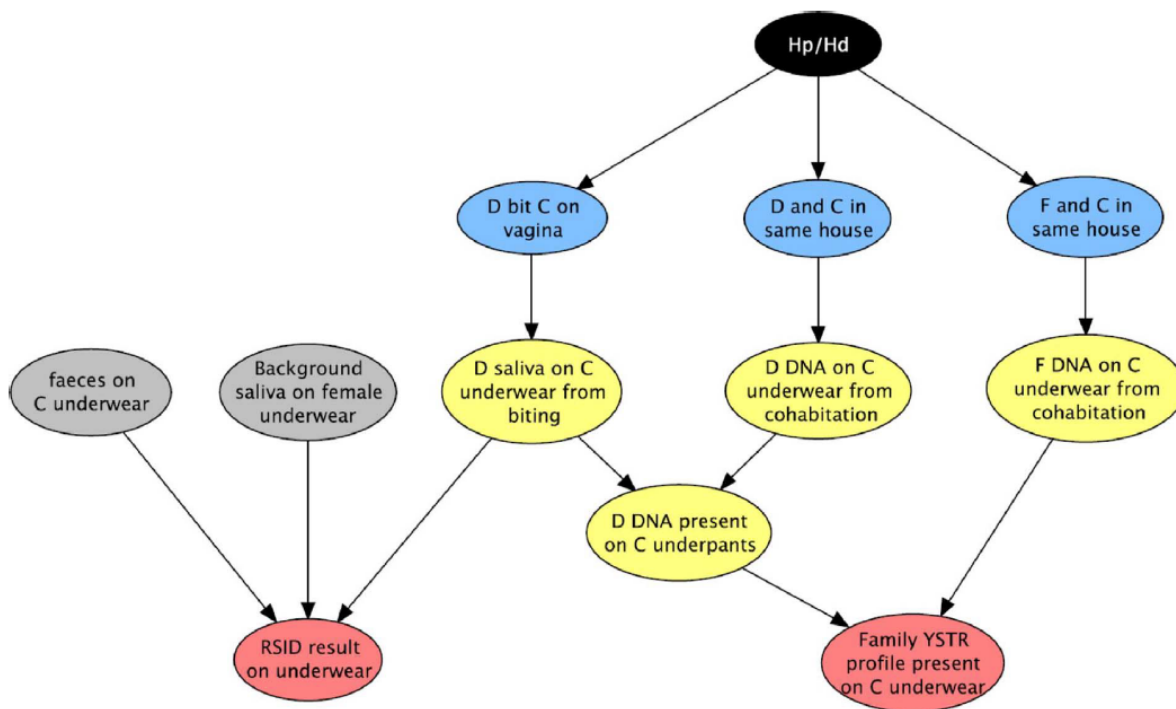


Figure 25: The BN taken from [47, p. 141].

The case modelled in Figure 25 involves an alleged assault of ‘D’ on ‘C’. This BN graph was constructed using the template by the same authors. The nodes are coloured to show which type of node they are: black for hypothesis node, blue for activity, yellow for transfer and persistence, red for findings, and grey for root nodes.

The activity-level forensic hypotheses identified in this case are:

H_p D has bitten C on the vagina, over her underwear.

H_d C has been staying at D’s home, but no biting occurred.

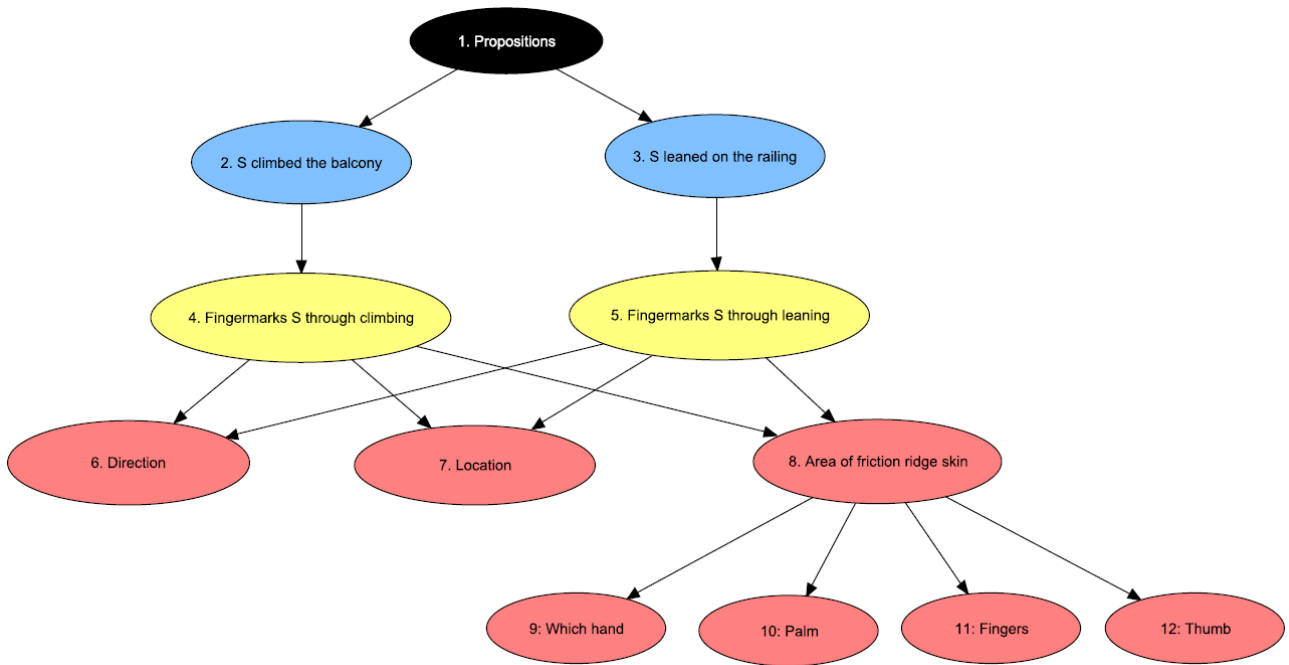


Figure 26: BN for the evaluation of fingerprints at activity level, taken from [42, p. 12]

The BN in Figure 26 models a burglary case with fingerprint evidence. Taylor and colleagues' method [47] was used in creating this BN.

The question in this case is what activity has led to the deposition of the fingerprints. The forensic hypotheses (at activity level) are:

H_p 'S' climbed the balcony and did not lean on the railing.

H_d 'S' leaned on the railing and did not climb the balcony.

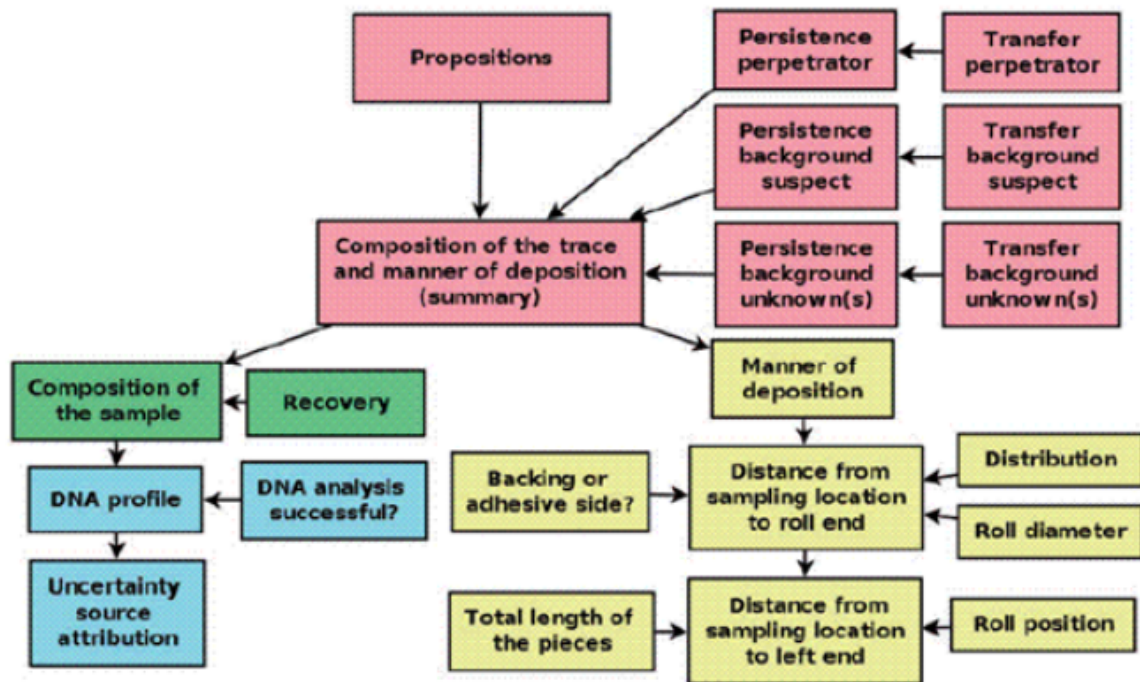


Figure 27: BN for the interpretation of single source biological traces on adhesive tape (taken from [55, p. 3])

This activity-level BN can be used to compute the probability of a suspect having taped an object given a human biological trace on a piece of adhesive tape.

The forensic hypotheses are:

H_p The suspect taped the object.

H_d An unknown person taped the object.

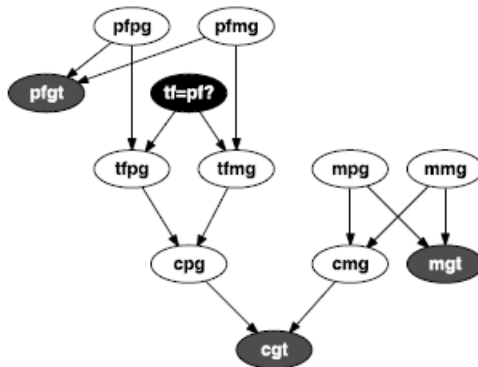


Figure 28: Simple paternity network, taken from [11, p. 580].

In the BN in Figure 28, there are nodes for each individual's genotype and paternal and maternal genes. For example, nodes 'cpg' and 'cmg' stand for the child's paternal and maternal gene, leading to the child's genotype in 'cgt'. The arcs reflect a definitional relationship.

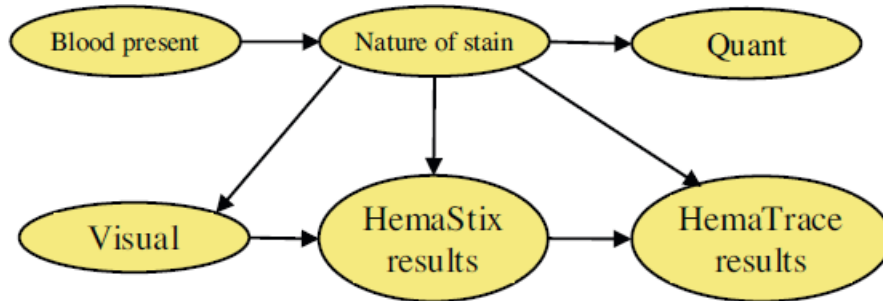


Figure 29: BN for evaluating the source of a biological trace (taken from [46, p. 56])

The BN in Figure 29 can be used to compute the probability of a victim being the donor of DNA recovered from a suspect's clothing. Two tests are carried out sequentially, a HemaStix and a HemaTrace test, and the arc from the former to the latter reflects the temporal relationship.

The source-level forensic hypotheses identified in this case are:

H_p The victim is the source of the DNA.

H_d SE (someone else) unrelated to the victim is the source of the DNA.

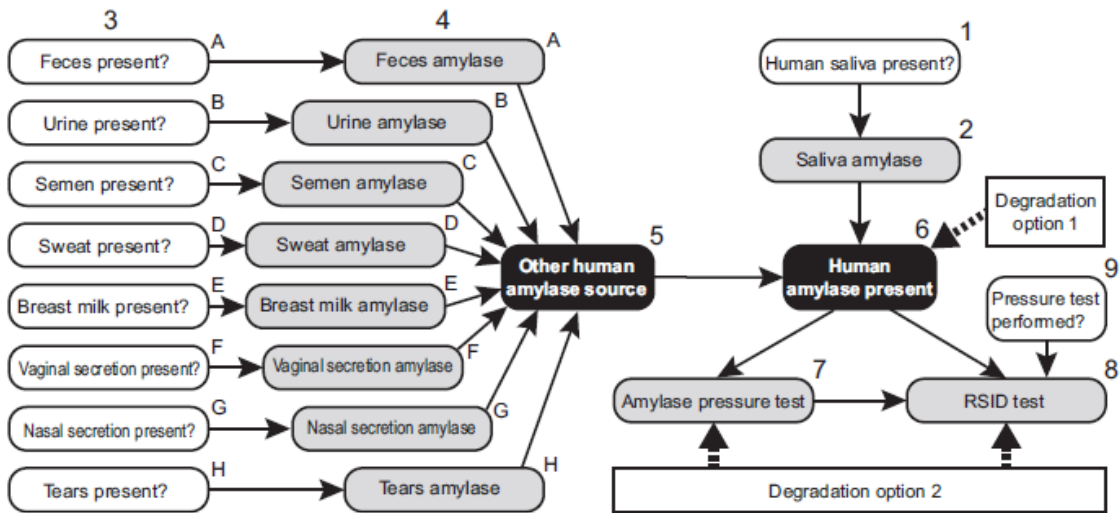


Figure 30: BN for saliva test results (taken from [57, p.328])

The BN in Figure 30 models the question of whether a forensic biology trace contains saliva or any other possible cause for the test results that is not saliva. The arc between nodes 7 and 8 was initially drawn during the manual construction, but subsequent checking against data revealed the influence of these nodes on each other was not as significant as the BN engineers thought.




Figure 31: BN for evaluating gunshot residue (GSR) particle evidence, taken from [4, p. 30]

Node 'D: condition of the stub' in Figure 31 leads to 'A: number of GSR particles present on lifting device prior to sampling' and in this way includes the uncertainty due to a contaminated sampling device.

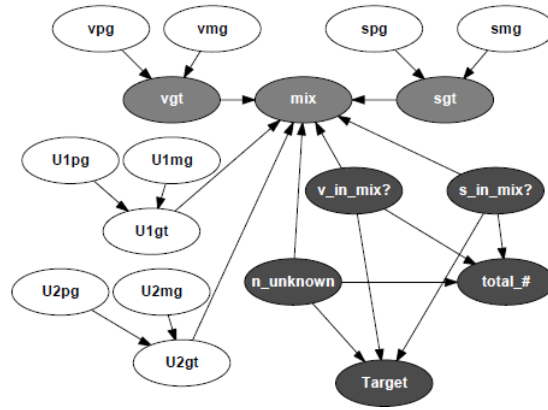


Figure 32: First attempt at modelling an unknown number of contributors for a mixture of DNA, taken from [34, p. 197]. Depending on the number of values established for the 'mix' node, its CPT contains at least 124416 entries.

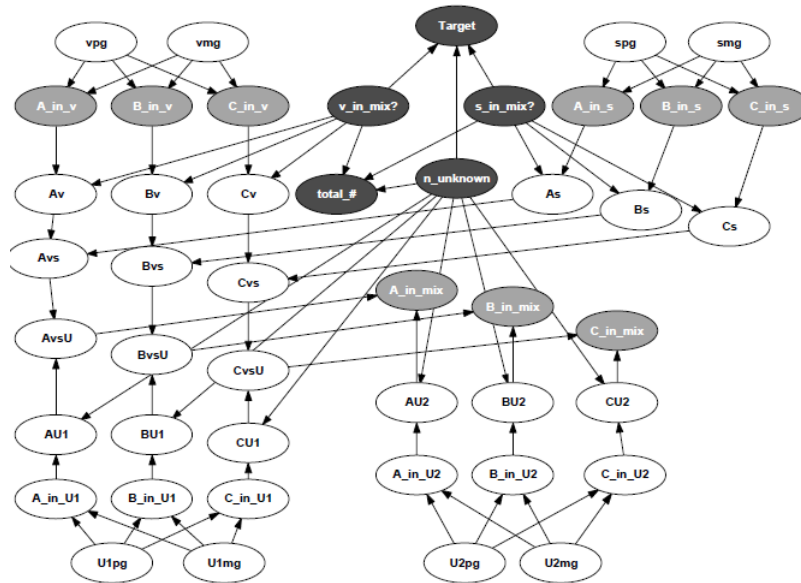


Figure 33: Second attempt at modelling an unknown number of contributors for a mixture of DNA, taken from [34, p. 198]

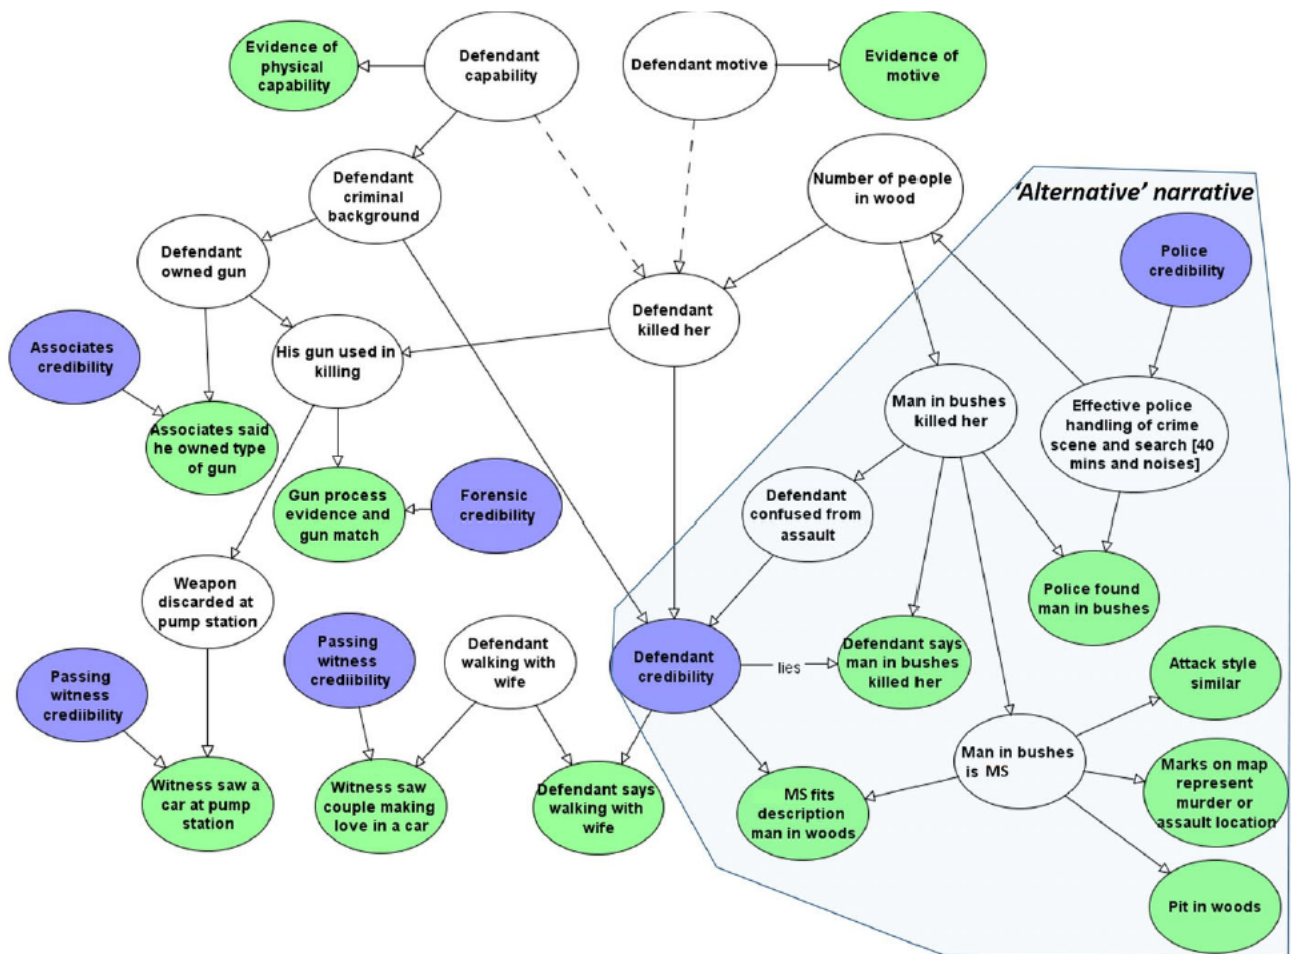


Figure 34: This BN models the murder case from the Simonshaven Woods, also known as the Simonshaven case, taken from [18, p. 11]

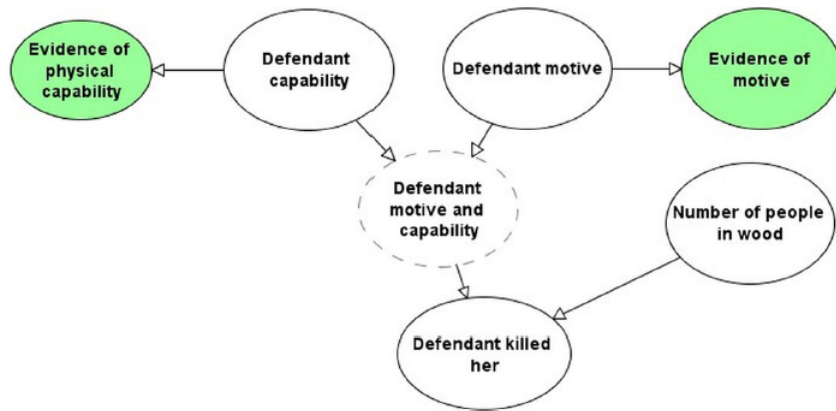


Figure 35: The dashed node represents the mediating AND-node capturing opportunity and motive from the Simonshaven BN. This node is hidden “underneath” dotted arcs in Figure 34. Taken from [18, p. 10]

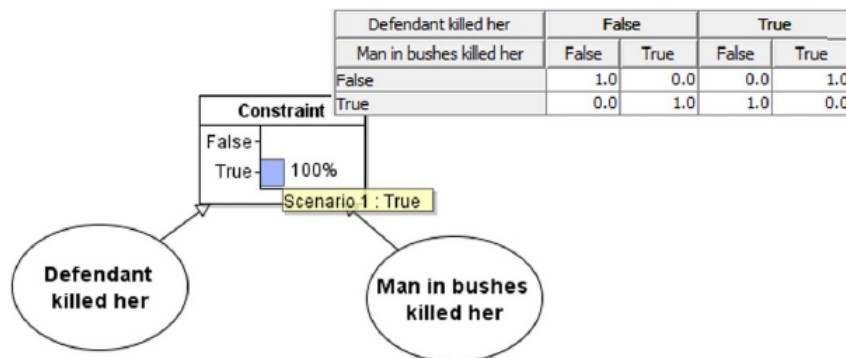


Figure 36: The ‘constraint’ node to enforce mutual exclusivity between the two hypothesis nodes in Fenton and colleagues Simonshaven BN from Figure 34, with its CPT. The constraint node is always set to ‘True’. This node is, like the mediating AND-node, hidden in the full BN. Taken from [18, p. 10]




Figure 37: BN modelling the probability of a pathological or toxicological cause of death (taken from [31, p. 7])

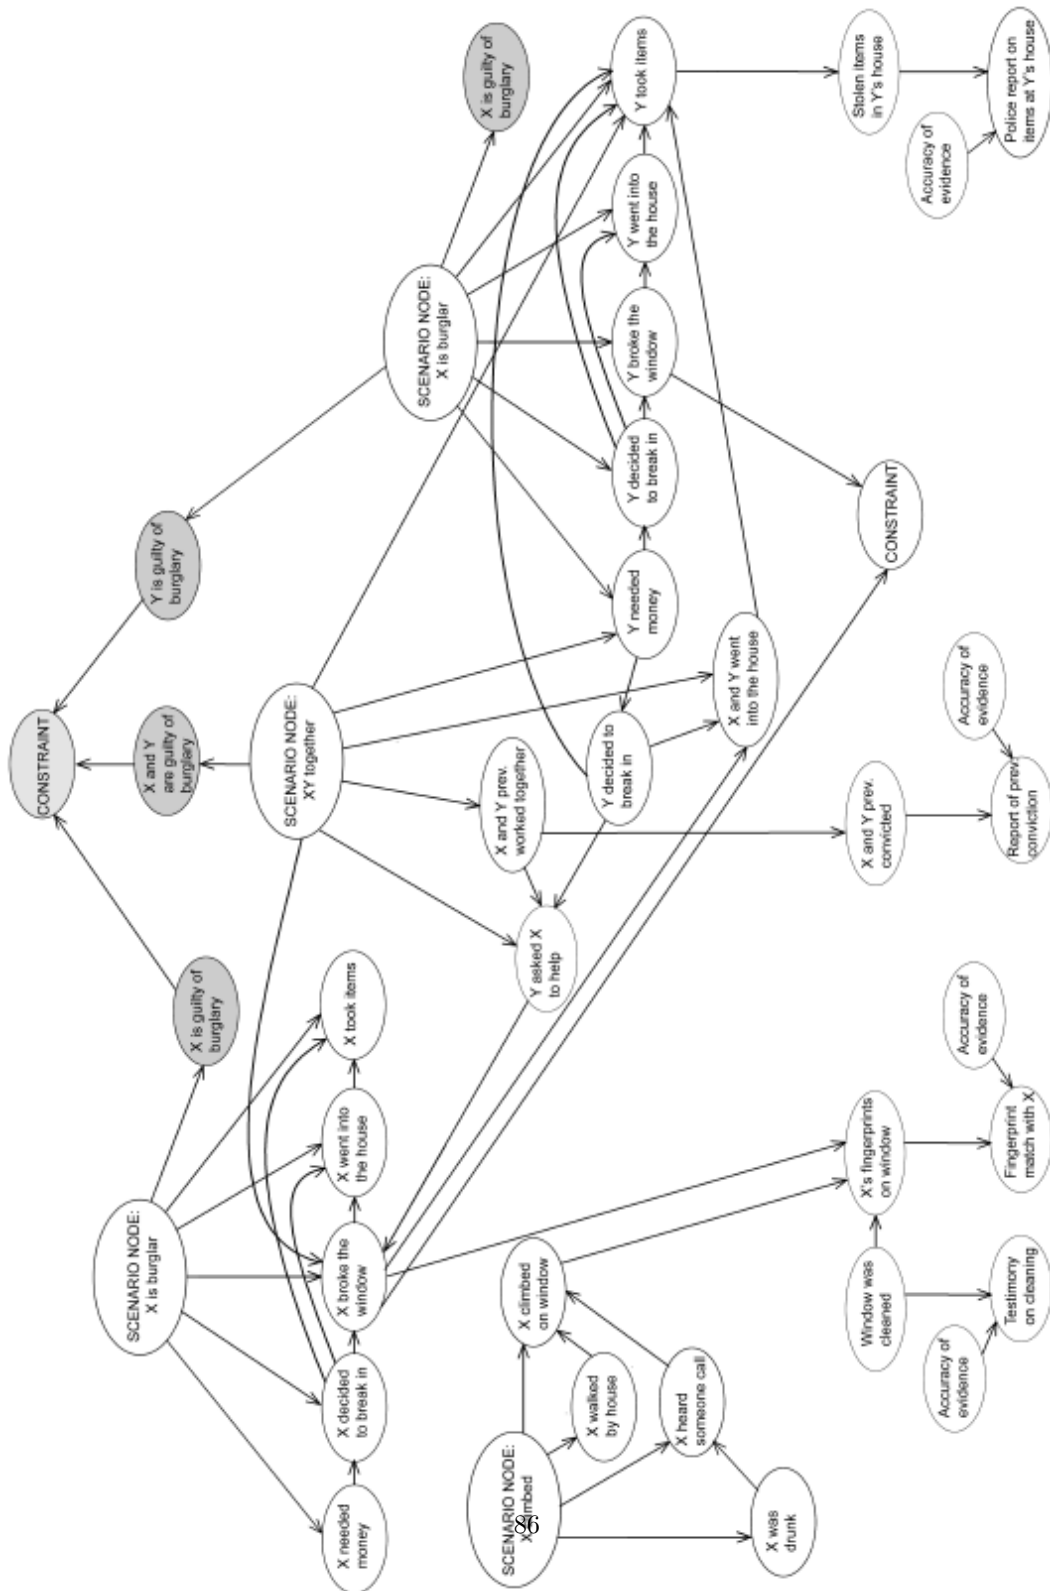


Figure 38: Burglar case, taken from [51, p. 9]. The BN includes a constraint node to ensure that mutually exclusive hypotheses cannot be true at the same time.

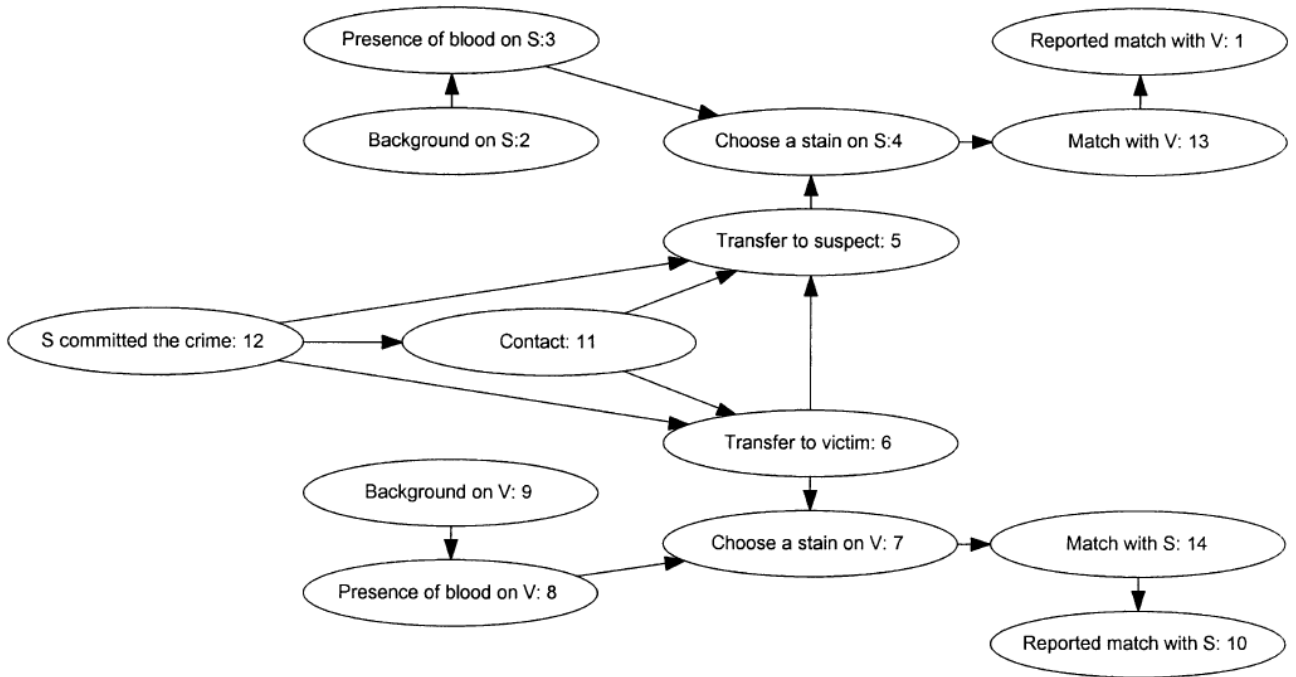


Figure 39: A BN for the evaluation of cross-transfer DNA evidence, taken from [1, p. 184]

B List of articles for the database of forensic and legal BNs

Author(s) and year	Forensic or legal	Level in hierarchy (if any)	Number of BNs collected from article
Aitken, Taroni, & Garbolino (2003)	Legal	Offence	1
Biedermann et al. (2005)	Forensic	Unspecified	1
Biedermann et al. (2005)	Forensic	Unspecified	2
Biedermann, Bozza & Taroni (2009)	Forensic	Unspecified	6
Dawid et al. (2002)	Forensic	Unspecified	5
Dawid (2003)	Forensic	Unspecified	1
Dawid, Mortera, & Vicard (2007)	Forensic	Unspecified	4
Evetts et al. (2001)	Legal & Forensic	Activity & Offence	2
Fenton, Neil & Lagnado (2012)	Legal	Offence	1
Fenton et al. (2019)	Legal	Offence	1
Haraksim et al. (2013)	Forensic	Source	2
Holický, Marková, & Sýkora (2013)	Forensic	Unspecified	1
Juchli, Biedermann & Taroni (2012)	Legal	Offence	1
Kokshoorn et al. (2017)	Forensic	Activity	4
Korb & Nicholson (2010)	Forensic	Unspecified	4
Kwan et al. (2008)	Forensic	Unspecified	1
Lagnado (2011)	Legal	Offence	1
Langford et al. (2015)	Forensic	Unspecified	1
McDermott & Aitken (2017)	Legal	Unspecified	1
Mortera et al. (2003)	Forensic	Unspecified	9
Oosterman et al. (2015)	Forensic	Source	1
De Ronde et al. (2019)	Forensic	Activity	3
Sironi et al. (2016)	Forensic	Unspecified	1
Smit et al. (2016)	Forensic	Source	1
Szkuta et al. (2018)	Forensic	Activity	3
Taylor et al. (2016)	Forensic	Source	3
Taylor et al. (2018)	Forensic	Activity	3
Vlek et al. (2013)	Legal	Unspecified	1
Wieten et al. (2015)	Forensic	Activity	1
De Wolff et al. (2015)	Forensic	Source	1
Zadora (2009)	Forensic	Source	1
De Zoete et al. (2016)	Forensic	Source	1

Table 4: List of articles from which forensic and legal BNs were collected to create a database. The full references can be found in the bibliography below.

Bibliography of database

This bibliography lists all articles from which BNs were collected to form a database of manually constructed BN graphs.

1. Aitken, C., Taroni, F., & Garbolino, P. (2003). A graphical model for the evaluation of cross-transfer evidence in DNA profiles. *Theoretical Population Biology*, 63(3), 179–190.
2. Biedermann, A., Taroni, F., Delemont, O., Semadeni, C., & Davison, A. (2005). The evaluation of evidence in the forensic investigation of fire incidents (Part I): An approach using Bayesian networks. *Forensic Science International*, 147(1), 49–57.
3. Biedermann, A., Taroni, F., Delemont, O., Semadeni, C., & Davison, A. C. (2005). The evaluation of evidence in the forensic investigation of fire incidents. Part II. Practical examples of the use of Bayesian networks. *Forensic Science International*, 147(1), 59–69.
4. Biedermann, A., Bozza, S., & Taroni, F. (2009). Probabilistic evidential assessment of gunshot residue particle evidence (Part I): Likelihood ratio calculation and case pre-assessment using Bayesian Networks. *Forensic Science International*, 191(1-3), 24–35.
5. Dawid, A. P., Mortera, J. Pascali, V. L., & Van Boxel, D. (2002). Probabilistic expert systems for forensic inference from genetic markers. *Scandinavian Journal of Statistics*, 29(4), 577–595.
6. Dawid, A. P. (2003). An object-oriented Bayesian network for estimating mutation rates. In *Proceedings of the Ninth International Workshop on Artificial Intelligence and Statistics*, Key West, Florida.
7. Dawid, A. P., Mortera, J., & Vicard, P. (2006). Representing and solving complex DNA identification cases using Bayesian networks. In A. Amorim & F. Corte-Real, N. Morling (Eds.), *Progress in Forensic Genetics 11, vol. 1288 of International Congress Series* (pp. 484–491).
8. Evett, I. W., Gill, P. D., Jackson, G., Whitaker, J., & Champod, C. (2001). Interpreting small quantities of DNA: The hierarchy of propositions and the use of Bayesian networks. *Journal of Forensic Science*, 47(3), 520–530.
9. Fenton, N., Neil, M., & D. Lagnado, D. A. (2012). A general structure for legal arguments about evidence using Bayesian networks. *Cognitive Science*, 37, 61–102.
10. Fenton, N., Neil, M., Yet, B., & Lagnado, D. (2019). Analyzing the Simonshaven case using Bayesian networks. *Topics in Cognitive Science*, 1–23.

11. Haraksim, R., Meuwly, D., Doekhie, G., Vergeer, P., & Sjerps, M. (2013). Assignment of the evidential value of a fingermark general pattern using a Bayesian network. In A. Brömme & C. Busch (Eds.), *BIOSIG 2013: Proceedings of the 12th International Conference of the Biometrics Special Interest Group* (pp. 99–103). Gesellschaft für Informatik.
12. Holický, M., Marková, J., & Sýkora, M. (2013). Forensic assessment of a bridge downfall using Bayesian networks. *Engineering Failure Analysis*, *30*, 1–9.
13. Juchli, P., Biedermann, A., & Taroni, F. (2012). Graphical probabilistic analysis of the combination of items of evidence. *Law, probability and risk*, *11*(1), 51–84.
14. Kokshoorn, B., Blankers, B. J., De Zoete, J., & Berger, C. E. (2017). Activity level DNA evidence evaluation: On propositions addressing the actor or the activity. *Forensic Science International*, *278*, 115–124.
15. Korb, K. B., & Nicholson, A. E. (2010). *Bayesian Artificial Intelligence* (2nd ed.). Chapman & Hall.
16. Kwan, M., Chow, K. P., Law, F., & Lai, P. (2008). Reasoning about evidence using Bayesian networks. In I. Ray & S. Shenoi (Eds.), *Advances in Digital Forensics IV* (pp. 275–289).
17. Lagnado, D. A. (2011). Thinking about evidence. In P. Dawid, W. Twining, & M. Vasilaki (Eds.), *Evidence, Inference and Enquiry. Proceedings of the British Academy*, *171* (pp. 183–223).
18. Langford, A. M., Bolton, J. R., Carlin, M. G., & Palmer, R. (2015). Post-mortem toxicology: A pilot study to evaluate the use of a Bayesian network to assess the likelihood of fatality. *Journal of Forensic and Legal Medicine*, *33*, 82–90.
19. McDermott, Y., & Aitken, C. (2017). Analysis of evidence in international criminal trials using Bayesian Belief Networks. *Law, Probability and Risk*, *16*(2-3), 111–129.
20. Mortera, J., Dawid, A. P., & Lauritzen, S. L. (2003). Probabilistic expert systems for DNA mixture profiling. *Theoretical Population Biology*, *63*(3), 191–205.
21. Oosterman, W. T., Kokshoorn, B., Maaskant-van Wijk, P. A., & De Zoete, J. (2015). From sub-source to source: Interpreting results of biological trace investigations using probabilistic models. *Forensic Science International: Genetics Supplement Series*, *5*, e590-e591.
22. De Ronde, A., Kokshoorn, B., De Poot, C. J., & De Puit, M. (2019). The evaluation of fingermarks given activity level propositions. *Forensic Science International*, *302*, [109904].

23. Sironi, E., Gallidabino, M., Weyermann, C., & Taroni, F. (2016). Probabilistic graphical models to deal with age estimation of living persons. *International Journal of Legal Medicine*, 130(2), 475–488.
24. Smit, N. M., Lagnado, D. A., Morgan, R. M., & Fenton, N. E. (2016). Using Bayesian networks to guide the assessment of new evidence in an appeal case. *Crime science*, 5(1), 1–12.
25. Szkuta, B., Ballantyne, K. N., Kokshoorn, B., & Van Oorschot, R. A. (2018). Transfer and persistence of non-self DNA on hands over time: Using empirical data to evaluate DNA evidence given activity level propositions. *Forensic Science International: Genetics*, 33, 84–97.
26. Taylor, D., Abaro, D., Hicks, T., & Champod, C. (2016). Evaluating forensic biology results given source level propositions. *Forensic Science International: Genetics*, 21, 54–67.
27. Taylor, D., Biedermann, A., Hicks, T., & Champod, C. (2018). A template for constructing Bayesian networks in forensic biology cases when considering activity level propositions. *Forensic Science International: Genetics*, 33, 136–146.
28. Vlek, C., Prakken, H., Renooij, S., Verheij, B. (2013). Modeling crime scenarios in a Bayesian network. In B. Verheij (Eds.), *Proceedings of the 14th International Conference on Artificial Intelligence and Law*, (pp. 150–159).
29. Wieten, R., De Zoete, J., Blankers, B., & Kokshoorn, B. (2015). The interpretation of traces found on adhesive tapes. *Law, Probability and Risk*, 14(4), 305–322.
30. De Wolff, T. R., Kal, A. J., Berger, C. E. H., & Kokshoorn, B. (2015). A probabilistic approach to body fluid typing interpretation: An exploratory study on forensic saliva testing. *Law, Probability and Risk*, 14(4), 323–339.
31. Zadora, G. (2009). Evaluation of evidence value of glass fragments by likelihood ratio and Bayesian Network approaches. *Analytica Chimica Acta*, 642(1-2), 279–290.
32. De Zoete, J., Oosterman, W., Kokshoorn, B., & Sjerps, M. (2016). Cell type determination and association with the DNA donor. *Forensic Science International: Genetics*, 25, 97–111.

Interview with forensic scientist

On June 19, 2020, we conducted a semi-structured interview with Bas Kokshoorn, a forensic DNA scientist at the Netherlands Forensic Institute (NFI).

The reasons behind interviewing this particular scientist were because he can be regarded both a BN engineer and domain expert, and because of his experience with manual BN construction. The main purpose of the interview was to gain insight from a BN engineer into the construction process of BNs at this department, in particular with respect to the BN graph. More specifically, we wanted to know what methods are used at this department at the NFI for BN (graph) construction, what motivations lie behind modelling choices and what aspects are considered in evaluating the graph. The interview took place online (via Skype) in the presence of the daily supervisor and lasted a little over two hours. Several days prior to the meeting, the interviewee was sent a brief summary of the literature discussed in Part I.

In this section we provide an outline of the interview in the form of a list of a summary of the main topics that were discussed:

- *Background information* The department of forensic biology/DNA traces at the NFI make daily use of BNs. These are constructed for forensic cases by order of the examining judge. In many cases BNs form the basis for the forensic scientists in formulating the LR, which are reported to the trier of fact. For simple cases no BNs are built, and the LR is simply computed from the formula (see [45]). The NFI provides an internal training on BN construction to its employees, and additional external courses can be taken at other forensic institutes.
- *BN graph construction method:* For construction of the graph, the forensic scientists at the department for forensic DNA investigation recently adopted the template by Taylor and colleagues (see Section 4). The steps from the method are not always explicitly carried out, but the rationale from the template is mainly followed: this method promotes the capacity of the graph to explain every possible route of transfer of DNA that is considered in the modelled case. Every activity is included in the BN, even non-disputed ones. The use of the same method across institutes ensures a degree of uniformity within the field and advances communication amongst BN engineers and domain experts.

In other departments BNs feature less strongly as a tool to assist the forensic practitioner, but with the recent increasing interest in BNs efforts are being made to develop a similar template. For other micro-traces the template might closely resemble Taylor and colleagues' template, but the template might not be so easily translated to other areas of forensic investigation. It does not provide any guidance on how to deal with other types of forensic traces that might involve more complex dependencies.

The idioms by Neil and colleagues (Section 3.2.3) are not explicitly used. Some parts of the BN graphs, such as the transfer of DNA for particular actions, always look the same. These could be viewed as 'idioms' or network fragments, but these are not documented as such. These 'idioms' are reused across different networks, sometimes with adjustments to the values (for example by switching from values representing absolute quan-

tities of DNA to relative quantities, or simply to the presence or absence of DNA).

Other techniques such as OOBNs are not used. Networks are generally not incorporated in other networks: a network for the evaluation of fingerprints, for instance, is not combined with one for DNA. The lack of data on the combination of variables from the networks prevents this, as no reliable probability estimations can be given.

- *Hierarchy of forensic hypotheses:* The hierarchy of forensic hypotheses (see Section 4.1) is used to evaluate the forensic biological evidence. In the early days the forensic hypotheses were often formed at source or sub-source level. With the advances in DNA technologies the source of a DNA trace is seldomly challenged, and the question shifted from source to activity.

For cases from departments other than human biology traces it can be difficult to establish the level from the hierarchy for the forensic hypotheses. These departments cannot benefit from the assistance the hierarchy provides in formulating the forensic hypotheses. In pathology for example, the main hypothesis mostly concerns the cause of death. A willing reader can view this as a source-level question, but this can also be seen as a matter of activities that resulted in death.

There is no difference in approach to BN construction for cases with forensic hypotheses at different levels. The vast majority of cases involves activity-level forensic hypotheses.

- *Desirable properties, adapted from Timmer:* In this part of the interview we asked about the concepts from the properties from Timmer (Section 5)²⁰ and whether the interviewee focused on these areas.

2 The explaining-away effect is not explicitly incorporated in the structure, or when determining the probability distributions (see [40]). For the field of DNA explaining-away is not always desired due to the fact that DNA can accumulate. To illustrate, one activity explaining the presence of a DNA trace does not automatically mean that another activity did not also happen. Other types of intercausal interactions are not considered either.

3-4 The computational feasibility of BNs is not an issue at the moment, as the structures are quite simplistic (even with the inclusion of every possible route of potential DNA transfer, i.e. inclusion of more nodes). Often additional nodes are included to provide insight into the considered transfer-routes and to ease the process of estimating probabilities, which means there are more parameters than strictly necessary. Divorcing is not used often, but occasionally mediating nodes (“summary nodes”) are included. Their function is rather to provide a summary

²⁰At the time of this interview, the new properties were not fully formed yet.

- 5 The translation of ‘premises’ and ‘conclusions’ to the forensic domain is not straightforward. A conversion of premises to ‘effects’ and conclusions to ‘causes’ is roughly what happens for the type of argumentation schemes considered by Timmer. Complications arise when considering ‘premises’ and ‘conclusions’ in this way for the pathology department for example, where the observed evidence are usually causes (of death). This property is not easily generalised.
- 6 Very occasionally BNs are used in case assessments to reason what direction to take in the forensic investigation. In this phase possible relevant scenarios are considered and BNs can be used to gain insight in the role of the evidence in differentiating these scenarios. Every once in a while a scenario, action, or piece of evidence is considered that is not connected to the hypothesis node and hence does not include an active chain between these nodes.
- 7 The direction of arcs as specified in property 7 (i.e. causally and temporally) follows implicitly from the template. The interviewee points out that this facilitates the parameter estimation. Ironically the reasoning in legal cases is the other way around, i.e. diagnostic: specific evidence is found, on the basis of which conclusions are drawn on the possible causes of this evidence.

- *Evaluation of the graph:* There is no readily available ‘checklist’ to evaluate the graphical characteristics of the BN. When a BN is used in the evaluation of evidence, the BN is always assessed by a second forensic scientist.

Constructional checks are done to ensure that the BN exhibits the desired behaviour. For two mutually exclusive nodes, the probabilities must be coordinated such that if one is true, the other must be false. Any arcs representing a conditional dependence which turns out to be non-existence upon evaluation of the parameters are removed. Doubts on the representation of conditional dependencies and independencies or the grouping of pieces of evidence are discussed with other colleagues within the department.

Initial BN graphs might be more complex in terms of the number of nodes and arcs than final versions. Sensitivity analyses help determining which nodes do not have a significant influence, such as background nodes in some cases.

- *Motivation behind modelling choices:* When starting out with the use of BNs at the NFI, their function was more a convenient way to organise the thoughts of the forensic scientist. Nowadays they are used to compute the LR. Thanks to the fact that the structure of BNs in this field is still relatively simple, the computational complexity is not an issue yet. With the increase in data on the probabilities of DNA transfer in particular situations, keeping the computational feasibility in check might be of relevance

in the future. The focus at the moment is ensuring the transparency of the graph with respect to communicating all possible manners of deposition of DNA. In other words, the emphasis is on the representational value of the graph.

No significant efforts are made to ensure mutual exclusivity between two mutually exclusive nodes (such as the ones described in the paragraph on property P1 in Section 7).

Author(s)	Year	Group (if applicable)	Topic	Specific BN (page nr)	Number of nodes	Number of arcs	Max(inc arcs)	Multiply connected ?	Nr h-t-h nodes	h-t-h/node ratio	Level in Hierarchy (Cook)	States	Reusable?
Atken et al	2003	Lausanne	cross-transfer evidence	184	14	16	3	yes	4	29%	Offence: 'S committed crime (or didn't)', note: mentioned that cross-transfer = activity	binary	yes: for cases with a single perp, a single suspect, a single victim and a single stain
Biedermann	2005a	Lausanne	Fire incidents pt I	53	8	7	2	No	2	25%	Unclear, cause of fire (natural, technical, human action) = proposition of interest.	binary, except H (ternary)	Yes, this BN is kind of a basic network for fire incidents and can be adjusted to specific cases (see Biedermann 2005b)
Biedermann	2005b	Lausanne	Fire incidents pt II	60-terpenes	8	7	2	no	3	38%	Unspecified	binary, except H (ternary)	case-specific
Biedermann	2005b	Lausanne	Fire incidents pt II	65-HPD	6	5	2	No	2	33%	Unspecified	non-binary, except M	case-specific
Biedermann	2009a	Lausanne	GSR simple	26	4	3	3	No	1	25%	Unspecified	binary, T and Y non-binary	Yes, basic structure that can be extended on
Biedermann	2009a	Lausanne	GSR alternative	27	3	2	2	No	1	33%	Unspecified	H binary, lambda and Y non-binary	Yes, basic structure that can be extended on, lambda can be defined according to needs of user
Biedermann	2009a	Lausanne	GSR + background	28	7	6	3	No	2	29%	Unspecified	all but T and S non-binary!	Yes, but is already an extension of previous
Biedermann	2009a	Lausanne	GSR + background + observed count	30	11	10	3	no	4	36%	Unspecified	all but T and S non-binary!	Yes, but is already an extension of previous
Biedermann	2009a	Lausanne	GSR + background + obs count + contamination	30	14	13	3	No	5	36%	Unspecified	all but T, D and S non-binary	Yes, but is already an extension of previous
Biedermann	2009a	Lausanne	GSR + background + obs count + contamination + case reconstruction	31	15	14	3	No	5	33%	Unspecified	all but T, D and S non-binary	Yes, but is already an extension of previous
Dawid	2002	QMUL	Forensic identification from DNA	280-simple	12	14	2	yes	7	58%	Unspecified	binary, ternary	yes
Dawid	2002	QMUL	Forensic identification from DNA	284-missing	18	22	2	yes	11	61%	Unspecified	binary, ternary	case-specific
Dawid	2002	QMUL	Forensic identification from DNA	285-missing	27	36	2	yes	18	67%	Unspecified	binary, ternary	case-specific
Dawid	2002	QMUL	Forensic identification from DNA	286- mutation simple	14	16	2	yes	7	50%	Unspecified	binary, ternary	yes
Dawid	2002	QMUL	Forensic identification from DNA	290-murder	43	56	2	yes	22	51%	Unspecified	binary, ternary	case-specific
De Ronde	2019	NFI	Fingermarks (activity I)	12-disp activity	12	14	2	yes	3	25%	Activity	uncertain	case-specific
De Ronde	2019	NFI	Fingermarks (activity I)	19-disp actor	17	27	4	yes	5	29%	Activity	uncertain	case-specific
De Ronde	2019	NFI	Fingermarks (activity I)	25->1 prints	34	66	5	yes	14	41%	Activity	uncertain	case-specific
de Wolff	2015	NFI	saliva	328	25	26	8	yes	4	16%	explicit activity, but maybe more source: what is the prob of human saliva being present (determine cell type, not whos or how)	binary, non-binary	yes
de Zoete	2016	NFI	cell type determination	2 (adapted from Onsterman)	15	17	4	yes	2	13%	source	binary, non-binary	yes
Evett	2002	-	DNA	523-burglary cigarette	5	4	4	No	1	20%	Activity	binary, except outcome (4)	case-specific
Evett	2002	-	DNA	527-watch	12	13	3	yes	3	25%	Offence	binary, 4	case-specific
Fenton	2012	QMUL	Legal	94	21	24	3	yes	8	38%	Offence	mostly binary	case-specific

	2019	QMUL	Simonshaven case	11	32	37	3	yes	10	31%	Offence	mostly binary, except number of people in wood	highly case-specific
Fenton													
Haraksim	2013	NFI	Fingermarks (source)	103-finger lvl	4	3	2	No	1	25%	Source	non-binary	yes
Haraksim	2013	NFI	Fingermarks (source)	106-person lvl	6	7	4	yes	1	17%	Source	non-binary	Yes
Holicky	2013	-	Bridge collapse	5	19	18	6	no	5	26%	Unspecified	binary	case-specific
Juchli	2012	Lausanne	finger & footwear mark	71	9	9	3	yes	3	33%	Offence	Mostly binary	case-specific
Kokshoorn	2017	NFI	Activity level DNA evidence evaluation	11	5	6	4	yes	1	20%	activity	binary, except DNA result (4)	yes Illustration of BN structures for specific sets of activity level propositions
Kokshoorn	2017	NFI	Activity level DNA evidence evaluation	13	5	5	4	yes	1	20%	activity	binary, except DNA result (4)	^^
Kokshoorn	2019	NFI	Activity level DNA evidence evaluation	15	5	4	4	No	0	0%	activity	binary, except DNA result (4)	^^
Kokshoorn	2019	NFI	Activity level DNA evidence evaluation	16	5	5	4	yes	1	20%	activity	binary, except DNA result (4)	^^
Korb	2010	Australia	Missing car	349-1	5	4	2	No	1	20%	activity?	binary	case-specific
Korb	2010	Australia	Missing car	349-2	5	5	2	yes	1	20%	activity?	binary	case-specific
Korb	2010	Australia	Missing car	349-3	4	3	1	No	0	0%	activity?	binary, ternary	case-specific
Korb	2010	Australia	Missing car	349-4 utility	6	5	2	No	1	17%	activity?	binary, ternary	case-specific
Kwan	2008	-	BitTorrent	285	24	28	3	yes	3	13%	Activity (could be offence)	ternary (y,n,uncertain)	case-specific
Kwan	2009	-	Yahoo!	249 (part)	5	4	1	no	0	0%	Activity	ternary (y,n,uncertain)	case-specific
Kwan, Overill	2011	-	Yahoo!	237 (full)	21	20	1	no	0	0%	Activity	ternary (y,n,uncertain)	case-specific
Lagnado	2011	QMUL	Legal case (Vole)	93	15	17	3	yes	5	33%	Offence	binary	case-specific
Langford	2015	-	Fatality (drugs)	7	10	9	8!	No	1	10%	Unspecified	binary	case-specific
McDermott	2017	-	criminal trials	14	7	11	2	yes	5	71%	Unspecified	binary	
Mortera	2003	QMUL, mix	DNA mixture	194-simple	16	20	2	yes	10	63%	Sub-source	binary, ternary	yes
Mortera	2003	QMUL, mix	DNA mixture	197-more contributors	24	30	3	yes	14	58%	Sub-source	binary, ternary	case-specific
Mortera	2003	QMUL, mix	DNA mixture	197-unknown n contribs (1)	18	21	7	yes	7	39%	Sub-source	binary, ternary	Yes
Mortera	2003	QMUL, mix	DNA mixture	198-unknown n contribs (2)	46	72	3	yes	35	76%	Sub-source	binary, ternary	Yes
Mortera	2003	QMUL, mix	DNA mixture	199-part	7	8	2	yes	4	57%	Sub-source	binary, ternary	Yes
Mortera	2003	QMUL, mix	DNA mixture	202-mixture + missing	22	28	2	yes	14	64%	Sub-source	binary, ternary	case-specific
Mortera	2003	QMUL, mix	DNA mixture	203-OJ simps. + silent	57	94	3	yes	46	81%	Sub-source	binary, ternary	case-specific
Mortera	2003	QMUL, mix	DNA mixture	204-mixture + missing+silent	25	31	2	yes	13	52%	Sub-source	binary, ternary	case-specific
Oosterman	2015	NFI	biological trace	e591	9	9	2	yes	2	22%	Source	uncertain	Yes, possibly extendable to more than 2 contributors, useable on its own as well as in another BN
Sironi													
Sironi	2015	Lausanne	Age estimation	479	3	2	1	No	0	0%	Unspecified	binary, non-binary	Yes
Smit	2016	QMUL	New evidence appeal case	6	7	8	2	yes	2	29%	source	binary	case-specific
Szkuta	2018	NFI, mix	secondary transfer of DNA	86	13	19	4	yes	6	46%	Activity	mostly binary	case-specific
Szkuta	2018	NFI, mix	secondary transfer of DNA + potential events	90	15	21	4	yes	6	40%	Activity	mostly binary	case-specific
Szkuta	2018	NFI, mix	secondary transfer of DNA + potential events + absence DNA	91	15	22	4	yes	7	47%	Activity	mostly binary	case-specific

Taylor	2016	Lausanne	presence of blood	56	6	7	2	yes	2	33%	Source	mostly binary, nature of stain (5, can be extended)	yes
Taylor	2016	Lausanne	include source of blood (single source)	60	8	9	2	yes	2	25%	Source	^^	yes
Taylor	2016	Lausanne	include mixture + lab error possibility	63	23	29	2	yes	9	39%	Source	^^	yes
Taylor	2018	Lausanne	DNA	141	12	13	3	yes	3	25%	Activity	mostly binary	case-specific
Taylor													
Taylor	2018	Lausanne	DNA: Alt Offender	144	20	24	4	yes	6	30%	Activity	mostly binary	case-specific
Vicard	2008	Mix	Mutation rates	13	8	9	3	yes	3	38%	Sub-source	binary, non-binary	Yes
Vlek	2013	UU	Burglary window	9	40	65	6	yes	25	63%	Activity	mostly binary	case-specific
Wieten	2015	NFI	biological traces on adhesive tapes	3	21	20	4	No	5	24%	Activity	binary, non-binary	yes
Wieten thesis	2014	NFI	biological traces on adhesive tapes	91	24	23	4	No	5	21%	Activity	binary, non-binary	yes
Wieten thesis	2014	NFI	biological traces on adhesive tapes: 2 single source traces	71	33	38	4	yes	11	33%	Activity	binary, non-binary	yes
Zadora	2009	-	Glass fragments	281	4	3	1	No	0	0%		binary, non-binary	yes
Zadora	2009	-	Glass fragments	284	10	12	3	yes	3	30%		binary, non-binary	yes

Author(s)	Year	motivation	Other remarks
Altken et al	2003	-	
Biedermann	2005a	-	
Biedermann	2005b	examine applicability of BNs to formal evidence analysis	
Biedermann	2005b	examine applicability of BNs to formal evidence analysis	
Biedermann	2009a	BN can handle events which would count as difficulties that complicate other computational procedures such as likelihood ratio formula: if one wished to incorporate other important factors such as suspect's activities during interval, this can be handled in graphical approach (esp. BN) where specifying variables and controlling additional variables is supported.	
Biedermann	2009a	AA	* Addition of network fragments to original model! Background presence * Size of P's CPT = large, 34221 entries! Filled with help of software
Biedermann	2009a		transfer+persistence = now 1 node
Biedermann	2009a		Node R is continuous but discretised
Biedermann	2009a		* Node A may be defined as a root node with probabilities of there being a = 0, 1, ... GSR particles, or alternatively as standard probab distribution (e.g. Poisson) * fragment D --> A K analogous to C --> B
Biedermann	2009a	"no claim of general applicability of any of models, mere purpose was emphasising flexibility and capacity of method".	
Dawid	2002	method can handle specific complications of forensic identification (on the basis of DNA), such as missing data on one or more individuals (samples from individuals of interest are not readily available and only indirectly relevant information is present), and genetic mutation (aka can handle uncertainty)	query node can be left out but preferred to be explicitly modelled due to simpler interpretation and direct access to the probability of interest aimed to reduce the computational burden by structuring the network at the most disaggregated level possible
Dawid	2002		
Dawid	2002		
Dawid	2002		
Dawid	2002		
De Ronde	2019	BNs in paper could function as basic networks for evaluation of fingermarks, extending/modifying to specific cases is possible, or function as building blocks. Makes process of evaluation of findings explicit, using network as tool for discussion of variables and dependencies	
De Ronde	2019		
De Ronde	2019		
de Wolff	2015	Explore possibilities	* arc introduced 7 --> 8: a positive result for 7 might predict a positive result for 8, but could also waste the sample: either way a dependency * later verified by data: suggests the dependency between 7 and 8 is not as strong as anticipated; arc may be obsolete (could be removed and then node 9 removed as well) --> however 7 + 8 remain to have an active chain, so no direct dependency but possibly not independent!
de Zoete	2016	present BN that can assist in associating donors and cell types (e.g. skin cells, semen, saliva)	
Evett	2002		Activity > X / SE is the person who smoked the cigarette
Evett	2002		Sub-source > DNA recovered from cigarette end came from suspect / unknown
Fenton	2012		

	2019	This approach is one in a series of analyses of the Simonshaven case; BN approach chosen to determine whether it was possible to construct one 'quickly' with a small group of people.	used building blocks (Fenton) opportunity, motive and intent != modelled explicitly, only relevant propositions are modelled Built with the idiom-based approach, idioms used: -Evidence accuracy -Cause-consequence -Opportunity & motive * Question regarding arc police credibility --> effective police handling, shouldn't it be police credibility --> police found man in bushes? In other 'testimonies' credibility node is linked directly * veracity, objectivity and competence are combined for simplicity into credibility
Fenton			
Haraksim	2013		
Haraksim	2013		
Hollicky	2013	take into account descriptions of expert judgments, illustrate variables and relationships, break down complex task into smaller sub-tasks for individual analysis, assess relative significance of individual causes	insufficiencies well documented by photographs, eyewitnesses, inspections, etc. --> this evidence not modelled (maybe not useful)
Juchli	2012		
Kokshoorn	2017		
Kokshoorn	2017		
Kokshoorn	2019		
Kokshoorn	2019		
Korb	2010	illustration	run into problem: $P(cm cs, sb)$: cs and sb = mutually exclusive events! 2 solutions: 2. And 3.
Korb	2010	illustration	
Korb	2010	illustration	
Korb	2010	illustration, decision node	
Kwan	2008		1 BN including 6 sub-hypotheses
Kwan	2009	used BN to quantify evidentiary strengths of hypotheses and to reason about evidence	
Kwan, Overill	2011	used BN to quantify evidentiary strengths of hypotheses and to reason about evidence	
Lagnado	2011		
Langford	2015	assist interpretation of post-mortem drug concentrations	
McDermott	2017		
Mortera	2003		
Mortera	2003		
Mortera	2003		max 2 unknown contributors (states total_# = {0, 1, 2}) + graphical representation = simple, intuitive - computationally inefficient!!!
Mortera	2003		
Mortera	2003	reformulation of subgraph of previous: allow use of simple arithmetic expressions to avoid tedious construction of states and CPTs for nodes	
Mortera	2003		
Mortera	2003		
Mortera	2003		
Oosterman	2015	BN can assist forensic practitioner in forming an opinion on cellular material (what) that sample is composed of and on attribution of donor (who) to cell type. Also: probabilistic approach can help improve quality of report (currently (dit was 2015..) no weight of evidence is given in report) and transparency of reasoning (association donor-cell type = complex, can lead to fallacies when reasoning)	
Sironi	2015		
Smit	2016	how forensic evidence should be presented in court	
Szkuta	2018		
Szkuta	2018		
Szkuta	2018		

Taylor	2016		* node "blood present" is not necessary to include, but serves as a summary (criterion 5) * arc HemaStix \exists HemaTrace: dependence: positive HemaStix means HemaTrace is more likely to be also positive. Could have been drawn other way around, but now temporal (mirrors testing in laboratory)
Taylor	2016		* 2 additional nodes are added to include sub-source info; Hp/Hd node has sub-source intermediate association propositions
Taylor	2016		* sub-sub-source propositions included (POI = source of major component of mixed DNA profile, etc.)
Taylor	2018	Paper aims to provide a method for constructing a BN on activity level propositions in biological trace cases	* Different choices in main proposition node formulation (disputed activity or disputed actor) will result in different BNs (vgl. Kokshoorn) * include undisputed activities (i.e. true in both Hp and Hd) - results in more dense BN + added complexity = necessary to properly account for non-disputed activities that may impact forensic results (B. van Tilburg mit 2.1)
Taylor	2018		For the previous BN ^ there was no need to consider background DNA on underwear as its ore/absence had no influence on PIE(H). Must now be included!
Vicard	2008	BN handles cases with incomplete information well	
Vlek	2013		
Wieten	2015	BN to aid forensic practitioners in interpretation of factors contributing to traces on adhesive tape	
Wieten thesis	2014	BN to aid forensic practitioners in interpretation of factors contributing to traces on adhesive tape	
Wieten thesis	2014		
Zadora	2009		
Zadora	2009		