# Probabilistic Reasoning 

(Probabilistisch Redeneren)

authors: Linda van der Gaag<br>Silja Renooij

Fall 2016

## Preface

In artificial-intelligence research, the Bayesian network framework for automated reasoning with uncertainty is has rapidly gained popularity since its introduction in the late 1980s. The framework provides a powerful formalism for representing a joint probability distribution on a set of statistical variables. In addition, it offers algorithms for efficient probabilistic inference. At present, more and more knowledge-based systems employing the framework are being developed for various domains of application, ranging from probabilistic information retrieval to medical diagnosis. This syllabus provides a tutorial introduction to the Bayesian network framework and highlights some issues of ongoing research in applying the framework for problem solving in real-life domains. Each chapter includes a number of exercises, and answers or hints to some of them (indicated by a ${ }^{*}$ ) are provided at the end.

This syllabus was first written in the late 1990s by L.C. van der Gaag and has been continuously under development ever since. Since 2001, adaptations and extensions have been made mostly by S. Renooij. The syllabus is by no means devoid from imperfections and any useful comments on its contents are greatly appreciated.

For the 2006 edition, several references to relevant recent research have been added to Chapters 4,5 , and 6 . In addition, material on the subject of sensitivity analysis has been extended. For the 2009 edition, some small errors were corrected and relevant references were added. Similar updates were done for the 2013 edition, for which also an example in Chapter 6 was extended. For the 2014 edition, an index was added, as well as relevant pointers to recent research. Notation was adapted to current standards for the 2016 edition.

Linda van der Gaag<br>Silja Renooij<br>Utrecht University<br>June 2016

© All rights reserved. No part of this work may be reproduced without permission of the authors.

## Contents

1 Introduction ..... 5
2 Preliminaries ..... 11
2.1 Graph Theory ..... 11
2.2 Probability Theory ..... 13
2.3 Notation ..... 17
3 Independences and Graphical Representations ..... 19
3.1 The Concept of Independence Revisited ..... 19
3.1.1 Pearl's Axiomatic System for Independence ..... 20
3.1.2 Properties of Independence Relations ..... 22
3.2 Graphical Representations of Independence ..... 23
3.2.1 Undirected Graphs ..... 24
3.2.2 Directed Graphs ..... 27
3.2.3 Choosing a Graphical Representation ..... 30
4 The Bayesian Network Framework ..... 37
4.1 The Bayesian Network Formalism ..... 37
4.2 Probabilistic Inference ..... 40
4.2.1 Directed Trees ..... 41
4.2.2 Singly Connected Digraphs ..... 54
4.2.3 Multiply Connected Digraphs ..... 62
4.2.4 Other Algorithms for Probabilistic Inference ..... 70
5 Building a Bayesian Network ..... 81
5.1 Identifying Variables and Values ..... 82
5.2 Constructing the Digraph ..... 84
5.2.1 Constructing the Digraph by Hand ..... 85
5.2.2 Learning the Digraph from Data ..... 85
5.3 Assessing Probabilities ..... 92
5.3.1 Sources of Probabilistic Information ..... 92
5.3.2 Simplifying Probability Assessment ..... 93
5.3.3 Eliciting Probabilities from Experts ..... 96
5.3.4 A Procedure for Probability Refinement ..... 97
6 Bringing Bayesian Networks into Practice ..... 105
6.1 Sensitivity Analysis ..... 106
6.1.1 What to Analyse? ..... 107
6.1.2 One-way Sensitivity Analysis ..... 108
6.1.3 Two-way Sensitivity Analysis ..... 114
6.2 Evaluating Bayesian Networks ..... 117
6.2.1 The Percentage Correct and its Shortcomings ..... 119
6.2.2 The evaluation score ..... 120
6.3 A Problem-Solving Architecture ..... 122
6.3.1 Threshold Decision Making ..... 123
6.3.2 Selective Evidence Gathering ..... 125
7 Conclusions ..... 131
8 Solutions, Answers and Hints ..... 133

## Chapter 1

## Introduction

Over the past few decades interest in the results of artificial-intelligence research has been growing to an increasing extent. Especially the area of intelligent or knowledgebased systems has attracted much attention. The phrase knowledge-based system is generally employed to denote computer systems in which some symbolic representation of human knowledge is incorporated and applied [Lucas \& Van der Gaag, 1991, Jackson, 1990]. Knowledge-based systems are typically designed to deal with real-life problems that require considerable human knowledge and expertise for their solution; in the early days they were therefore called expert systems. Examples of such systems range from medical diagnosis and technical trouble shooting to financial advice and product design. It is their ability to capture and reason with (specialised) human knowledge that allows knowledge-based systems to arrive at a performance comparable to that of human experts. Knowledge-based systems by now have found their way from academic laboratories to the industrial world and are being integrated into conventional software environments.

As more and more knowledge-based systems are being developed for a large variety of problems, it becomes apparent that the knowledge required to solve these problems often is not precisely defined but instead is of an imprecise nature. In fact, many reallife problem domains are fraught with uncertainty. Human experts in these domains typically are able to form judgements and take decisions based on uncertain, incomplete, and sometimes even contradictory information. To be of practical use, a knowledgebased system has to deal with such information at least equally well. For this purpose, a knowledge-based system employs a formalism for representing uncertainty and an associated algorithm for manipulating uncertain information. The major research topic in artificial intelligence of reasoning with uncertainty, or plausible reasoning, addresses the design of such formalisms and algorithms [Shafer \& Pearl, 1990].

As probability theory is a mathematically well-founded theory about uncertainty, having a long and outstanding tradition of research and experience, it is not surprising that this theory takes a prominent place in research on reasoning with uncertainty in knowledge-based systems. Unfortunately, applying probability theory in a knowledgebased context is not as easy as it may seem at first sight. Straightforward application of the basic concepts from probability theory leads to insuperable problems of computational complexity: explicit representation of a joint probability distribution requires exponential space (exponential in the number of variables discerned), and even if the distribution could be represented more economically, computing probabilities of interest by the basic rules of marginalisation and conditioning would have an exponential
time complexity. The rich history of applying probability theory for reasoning with uncertainty in knowledge-based systems shows various attempts to settle these problems.

In this chapter, we sketch the historical background of applying probability theory in a knowledge-based system. We would like to note that our intention is not to be complete, but merely to give an impression of the problems encountered by researchers pioneering in automated probabilistic inference. In our sketch, we focus on the task of (medical) diagnosis. For a given problem domain, we discern a set of possible hypotheses $\boldsymbol{H}=\left\{h_{1}, \ldots, h_{n}\right\}, n \geq 1$, and a set of pieces of evidence $\boldsymbol{E}=\left\{e_{1}, \ldots, e_{m}\right\}, m \geq 1$, that may be observed in relation with these hypotheses. For ease of exposition, we assume that each of the hypotheses is either true or false; equally, we assume that each of the pieces of evidence is either true or false. A diagnostic problem in this domain now is a set of pieces of evidence $\boldsymbol{e} \subseteq \boldsymbol{E}$ that is actually observed and needs to be explained in terms of the hypotheses discerned. A diagnosis for a problem $\boldsymbol{e}$ under consideration is a set of hypotheses $\boldsymbol{h} \subseteq \boldsymbol{H}$ that best explains $\boldsymbol{e}$.

As early as in the 1960s several research efforts on automated reasoning with uncertainty for diagnostic applications were undertaken [Warner et al., 1961, Gorry \& Barnett, 1968, De Dombal et al., 1972]. The systems constructed in this period were based to a large extent on application of Bayes' Theorem; in the sequel, we will refer to the approach taken in these early systems as the naive-Bayesian approach. In this approach, the basic idea of computing a diagnosis for a set of actually observed pieces of evidence $\boldsymbol{e} \subseteq \boldsymbol{E}$ is to compute for all sets of hypotheses $\boldsymbol{h} \subseteq \boldsymbol{H}$ the conditional probability $\operatorname{Pr}(\boldsymbol{h} \mid \boldsymbol{e})$ from the distribution $\operatorname{Pr}$ on the domain at hand, and then select a set $\boldsymbol{h}^{\prime} \subseteq \boldsymbol{H}$ with highest probability. Since for real-life applications the conditional probabilities $\operatorname{Pr}(\boldsymbol{e} \mid \boldsymbol{h})$ often are easier to come by than the conditional probabilities $\operatorname{Pr}(\boldsymbol{h} \mid \boldsymbol{e})$, generally Bayes' Theorem is used for computing the required probabilities:

$$
\operatorname{Pr}(\boldsymbol{h} \mid \boldsymbol{e})=\frac{\operatorname{Pr}(\boldsymbol{e} \mid \boldsymbol{h}) \cdot \operatorname{Pr}(\boldsymbol{h})}{\operatorname{Pr}(\boldsymbol{e})}
$$

It will be evident that this approach is quite expensive from a computational point of view: because a diagnosis may be composed of several different hypotheses, the number of probabilities to be computed equals $2^{n}-1$. To overcome this problem of time complexity, a simplifying assumption is made: it is assumed that all hypotheses are mutually exclusive and collectively exhaustive. With this assumption only the $n$ singleton hypothesis sets $\left\{h_{i}\right\}$ have to be considered as possible diagnoses. So, only the probabilities $\operatorname{Pr}\left(h_{i} \mid \boldsymbol{e}\right)$ (writing $h_{i}$ instead of $\left\{h_{i}\right\}$ ) for all $h_{i} \in \boldsymbol{H}$ have to be computed. To this end, once more Bayes' Theorem is used:

$$
\operatorname{Pr}\left(h_{i} \mid \boldsymbol{e}\right)=\frac{\operatorname{Pr}\left(\boldsymbol{e} \mid h_{i}\right) \cdot \operatorname{Pr}\left(h_{i}\right)}{\operatorname{Pr}(\boldsymbol{e})}=\frac{\operatorname{Pr}\left(\boldsymbol{e} \mid h_{i}\right) \cdot \operatorname{Pr}\left(h_{i}\right)}{\sum_{k=1}^{n} \operatorname{Pr}\left(\boldsymbol{e} \mid h_{k}\right) \cdot \operatorname{Pr}\left(h_{k}\right)}
$$

For automated application of Bayes' Theorem in this form, several probabilities are required from the joint probability distribution $\operatorname{Pr}$ on the domain at hand. In fact, conditional probabilities $\operatorname{Pr}\left(\boldsymbol{e} \mid h_{k}\right), k=1, \ldots, n$, for every combination of pieces of evidence $\boldsymbol{e} \subseteq \boldsymbol{E}$, have to be available. Apart from the fact that it is hardly likely that these probabilities will be readily available in a real-life problem domain, this means storing exponentially many probabilities. To overcome this problem of space complexity, a second simplifying assumption is made: it is assumed that all pieces of evidence are conditionally independent given any of the hypotheses discerned.

The two simplifying assumptions taken together allow for computing the probabilities $\operatorname{Pr}\left(h_{i} \mid \boldsymbol{e}\right)$ for all $h_{i} \in \boldsymbol{H}$ given observed evidence $\boldsymbol{e}=\left\{e_{j_{1}}, \ldots, e_{j_{p}}\right\}, 1 \leq p \leq m$, from

$$
\begin{aligned}
\operatorname{Pr}\left(h_{i} \mid e_{j_{1}} \wedge \cdots \wedge e_{j_{p}}\right) & =\frac{\operatorname{Pr}\left(e_{j_{1}} \wedge \cdots \wedge e_{j_{p}} \mid h_{i}\right) \cdot \operatorname{Pr}\left(h_{i}\right)}{\sum_{k=1}^{n} \operatorname{Pr}\left(e_{j_{1}} \wedge \cdots \wedge e_{j_{p}} \mid h_{k}\right) \cdot \operatorname{Pr}\left(h_{k}\right)}= \\
& =\frac{\operatorname{Pr}\left(e_{j_{1}} \mid h_{i}\right) \cdots \operatorname{Pr}\left(e_{j_{p}} \mid h_{i}\right) \cdot \operatorname{Pr}\left(h_{i}\right)}{\sum_{k=1}^{n} \operatorname{Pr}\left(e_{j_{1}} \mid h_{k}\right) \cdots \operatorname{Pr}\left(e_{j_{p}} \mid h_{k}\right) \cdot \operatorname{Pr}\left(h_{k}\right)}
\end{aligned}
$$

It will be evident that for any diagnostic problem $\boldsymbol{e}$ now only $n-1$ probabilities have to be computed, and that for this purpose only $m \cdot n$ conditional probabilities and $n-1$ prior ones have to be stored.

The systems for automated reasoning with uncertainty constructed in the 1960s were rather small-scaled: they were devised for clear-cut problem domains with only a small number of hypotheses and restricted evidence. For these small systems, all probabilities necessary for applying Bayes' Theorem could be acquired from statistical analysis of empirical data. Despite the underlying (over-)simplifying assumptions, these systems performed considerably well [De Dombal et al., 1974]. Nevertheless, interest in this naive Bayesian approach to reasoning with uncertainty faded in the late 1960s and early 1970s. One of the reasons for this decline in interest is that the approach was feasible only for highly restricted problem domains. For larger or more complex domains, the above-mentioned simplifying assumptions often were seriously violated, causing degeneration of system behaviour. In addition, for larger domains the approach inevitably became demanding, either computationally or from an assessment point of view.

At this stage, the first diagnostic knowledge-based systems began to emerge from artificial-intelligence research. These systems mostly use production rules for representing human (experiential) knowledge in a modular form closely resembling logical implications - production rules are expressions of the form if $\langle$ condition $\rangle$ then $\langle$ conclusion $\rangle$. These so-called rule-based expert systems exhibit 'intelligent' reasoning behaviour by employing a heuristic reasoning algorithm that use the production rules for selective gathering of evidence and for pruning the search space of possible diagnoses. It is this pruning behaviour that renders the rule-based expert systems capable of dealing with larger and complexer problem domains than the early naive-Bayesian systems are. The best-known rule-based expert system developed in the 1970s is the MYCIN system for assisting physicians in the diagnosis and treatment of bacterial infections [Buchanan \& Shortliffe, 1984].

In the context of rule-based expert systems, the naive Bayesian approach to reasoning with uncertainty is no longer feasible due to the large number of probabilities to be computed. Since in a rule-based system during problem solving the search space of possible diagnoses is pruned by heuristic as well as probabilistic criteria, it is necessary to compute probabilities for all intermediate results derived by the production rules in addition to the probabilities of the separate hypotheses. To allow for efficient computation of all these probabilities, a set of computation rules has been designed. These computation rules provide for computing the probability of an (intermediate) result from probabilities associated with the production rules that are used in its derivation; to this end, each production rule is assigned the conditional probability of its conclusion
given its condition. Unfortunately, these computation rules do not always accord with the axioms of probability theory and can not even be considered approximation rules for computing probabilities. In the sequel, we will use the phrase quasi-probabilistic to refer to this approach. The most well-known illustration of the quasi-probabilistic approach is the certainty-factor model, designed originally for dealing with uncertainty in the MYCIN system [Shortliffe \& Buchanan, 1984]. The certainty-factor model enjoys widespread use in rule-based expert systems built after MYCIN, even though by now it is widely known that the model is mathematically flawed. The relative success of the model can however be accounted for by its satisfactory behaviour in most applications and by its conceptual and computational simplicity [Van der Gaag, 1994].

Although the quasi-probabilistic approach to reasoning with uncertainty in knowl-edge-based systems on the one hand met with considerable success in the artificialintelligence community, it was criticised severely on the other hand because of its adhoc character. The incorrectness of the approach from a mathematical point of view even led to a world-wide debate concerning the appropriateness of probability theory for handling uncertainty in a knowledge-based context.

The adversaries of probability theory argue that the theory is not expressive enough to cope with the different kinds of uncertainty that are encountered in real-life problem domains and therefore have to be dealt with in knowledge-based systems. As a consequence several other (more or less) mathematical models have been proposed for reasoning with uncertainty. A major trend in plausible reasoning has arisen from the claim that probability theory is not able to capture imprecision or vagueness, notions of uncertainty which are salient in natural language representations. The name of L.A. Zadeh is inseparable from this trend: he was the first to propose fuzzy set theory as the point of departure for the development of methods that are able to cope with vague information. Dempster-Shafer theory lies at the basis of another major trend in plausible reasoning. The theory was developed by G. Shafer, building on earlier work by A.P. Dempster [Shafer, 1976]. It was motivated by the observation that probability theory is not able to discern between uncertainty and ignorance due to incompleteness of information.

The advocates of probability theory, on the other hand, claim that it is provable that probability theory is the only correct way of dealing with uncertainty and that anything that can be done with non-probabilistic methods, can be done equally well using a probability-based method. For this claim often an argument by R.T. Cox is cited [Cox, 1979]: Cox states a simple set of intuitive properties a measure of uncertainty has to satisfy and subsequently shows that the basic axioms of probability theory follow. Here, we will not enter into the debate concerning the appropriateness of probability theory for reasoning with uncertainty in knowledge-based systems; for a wide range of diverging opinions, the reader is referred to [Cheeseman, 1988] with its ensuing discussions.

Although the above-mentioned debate was not in the least subdued, in the mid1980s the probabilistic network framework was introduced as a novel approach to applying probability theory for reasoning with uncertainty in knowledge-based systems [Pearl, 1988]. The probabilistic network framework is characterised by a powerful formalism for representing domain knowledge and the uncertainties that go with it - more in specific, the formalism provides for a concise representation of a joint probability distribution on a set of statistical variables. Associated with this formalism are algorithms for efficiently computing probabilities of interest and for processing evidence; these al-
gorithms constitute the basic building blocks for reasoning with knowledge represented in the formalism. When compared to the naive-Bayesian approach on the one hand and the quasi-probabilistic approach on the other hand, the probabilistic network approach offers advantages over both. In contrast with the quasi-probabilistic approach, the probabilistic network approach has a firm mathematical foundation in probability theory. Contrasting the naive-Bayesian approach, the probabilistic network approach circumvents the need for simplifying assumptions by capturing and reasoning about actual independences among variables.

Various names have been used to refer to probabilistic networks, such as Bayesian belief networks, belief networks, causal networks and Bayesian networks. Nowadays, the general family of graphical-model-based probabilistic networks is known by the phrase Probabilistic Graphical models (PGMs); the discrete and directed family member which is the topic of this syllabus is mostly referred to as Bayesian Network. Since its introduction, the Bayesian network framework has rapidly gained in popularity and by now is beginning to illustrate its worth in complex problem domains: practical applications have been and are being developed for example for medical diagnosis and prognosis [Andreassen et al., 1987, Heckerman et al., 1992, Blanco et al., 2005], for probabilistic information retrieval [Bruza \& Van der Gaag, 1994], in computer vision [Jensen et al., 1990], in forensic science [Taroni et al., 2006] and various other domains (see e.g. [Pourret, Naim \& Marcot, 2008]). Whereas earlier applications of Bayesian networks were mostly handcrafted with the help of domain experts, the increasing availability of large data sets has made it much easier to construct applications directly from data [Neapolitan, 2003]. Even large data sets, however, usually do not contain sufficient reliable information to construct reliable networks of general topology; for this reason, network engineers either resort to using various types of classifier [Friedman et al., 1997], or again have to rely on domain expertise to complete the network [Druzdzel \& Van der Gaag, 2000].

This syllabus provides a tutorial introduction to the Bayesian network framework and highlights some issues of ongoing research in applying the framework for real-life problem solving. It is organised as follows. Chapter 2 provides some preliminaries from graph theory and from probability theory. In Chapter 3, we discuss the representation of probabilistic independence in graphical models. Chapter 4 introduces the Bayesian network framework: it details the Bayesian network formalism and outlines its associated algorithms. In Chapter 5 we address building Bayesian networks for real-life problem domains. Analysis of and problem solving with Bayesian networks is the topic of Chapter 6. The syllabus is rounded off with some conclusive discussion in Chapter 7.

## Chapter 2

## Preliminaries

### 2.1 Graph Theory

In this section, some concepts from graph theory are reviewed. Our review is tailored to the Bayesian network framework and is not meant to be exhaustive; for further information, any introductory textbook on graph theory will suffice.

Generally two types of graph are discerned: undirected and directed ones.
Definition 2.1.1 An undirected graph $G$ is a pair $G=\left(\boldsymbol{V}_{G}, \boldsymbol{E}_{G}\right)$ where $\boldsymbol{V}_{G}$ is a finite set of vertices (also called nodes) and $\boldsymbol{E}_{G}$ is a set of unordered pairs $\left(V_{i}, V_{j}\right), V_{i}, V_{j} \in$ $\boldsymbol{V}_{G}$, called edges.

A directed graph, or digraph for short, is a pair $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ where $\boldsymbol{V}_{G}$ is a finite set of vertices and $\boldsymbol{A}_{G}$ is a set of ordered pairs $\left(V_{i}, V_{j}\right), V_{i}, V_{j} \in \boldsymbol{V}_{G}$, called arcs. An arc $\left(V_{i}, V_{j}\right)$ is often written as $V_{i} \rightarrow V_{j}$ or $V_{j} \leftarrow V_{i}$.

For a vertex in a graph, different sets of related vertices can be identified.
Definition 2.1.2 In a digraph $G$, vertex $V_{j}$ is called a predecessor (or parent) of vertex $V_{i}$ if $\left(V_{j}, V_{i}\right) \in \boldsymbol{A}_{G}$; the set of all predecessors of vertex $V_{i}$ in $G$ is denoted as $\boldsymbol{\rho}_{G}\left(V_{i}\right)$. Likewise, vertex $V_{j}$ is called a successor (or child) of vertex $V_{i}$ if $\left(V_{i}, V_{j}\right) \in \boldsymbol{A}_{G}$; the set of all successors of vertex $V_{i}$ in $G$ is denoted as $\boldsymbol{\sigma}_{G}\left(V_{i}\right)$. The reflexive transitive closure ${ }^{1}$ of $V_{i}$ under the predecessor relation is denoted as $\rho_{G}^{*}\left(V_{i}\right)$; an element from $\boldsymbol{\rho}_{G}^{*}\left(V_{i}\right)$ is called an ancestor of $V_{i}$. Similarly, $\boldsymbol{\sigma}_{G}^{*}\left(V_{i}\right)$ denotes the descendants of $V_{i}$.
The set of neighbours of vertex $V_{i}$ is defined as

$$
\boldsymbol{\nu}_{G}\left(V_{i}\right)= \begin{cases}\boldsymbol{\sigma}_{G}\left(V_{i}\right) \cup \boldsymbol{\rho}_{G}\left(V_{i}\right) & \text { if } G \text { is directed; } \\ \left\{V_{j} \mid\left(V_{i}, V_{j}\right) \in \boldsymbol{E}_{G}\right\} & \text { if } G \text { is undirected }\end{cases}
$$

The size of the neighbour-set of a vertex is called its degree. In case of a vertex in a digraph, we in addition define the in-degree to be its number of predecessors and the out-degree to be the number of its successors; the incoming and outgoing arcs together are called its incident arcs.

In the sequel, we will often drop the subscript $G$ from $\boldsymbol{\rho}_{G}$ etc. as long as ambiguity cannot occur.

[^0]The following definition introduces several types of vertex sequence for undirected graphs.

Definition 2.1.3 Let $G=\left(\boldsymbol{V}_{G}, \boldsymbol{E}_{G}\right)$ be an undirected graph. A path from vertex $V_{0}$ to vertex $V_{k}$ in $G$ is a sequence of vertices $V_{0}, \ldots, V_{k}, k \geq 0$, with distinct edges $\left(V_{i-1}, V_{i}\right) \in$ $\boldsymbol{E}_{G}, i=1, \ldots, k$, between them; $k$ is called the length of the path. A path is called simple if all vertices are distinct. A cycle is a path $V_{0}, \ldots, V_{k}, V_{0}$ from $V_{0}$ to $V_{0}$ of non-zero length. The graph $G$ is called cyclic if it contains at least one cycle; otherwise, it is called acyclic.

In undirected graphs, self-loops (an edge $\left(V_{0}, V_{0}\right)$ ) are generally not allowed. The concepts of path and cycle introduced for undirected graphs directly apply to directed graphs by considering arcs rather than edges. Unless stated otherwise, we typically assume paths to be simple.

We now introduce the concept of underlying graph. This concept associates an undirected graph with a directed one. We thereby assume that directed graphs do not contain self-loops either, although this is not a convention.

Definition 2.1.4 Let $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ be a digraph. The underlying graph $H$ of $G$ is the undirected graph $H=\left(\boldsymbol{V}_{H}, \boldsymbol{E}_{H}\right)$ where $\boldsymbol{V}_{H}=\boldsymbol{V}_{G}$ and $\boldsymbol{E}_{H}$ is obtained from $\boldsymbol{A}_{G}$ by replacing each arc $\left(V_{i}, V_{j}\right) \in \boldsymbol{A}_{G}$ by an edge $\left(V_{i}, V_{j}\right)$.

Related to a digraph's underlying graph, we introduce two additional types of vertex sequence for digraphs.

Definition 2.1.5 Let $G$ be a digraph and let $H$ be its underlying graph. A chain from vertex $V_{0}$ to vertex $V_{k}$ in $G$ is a sequence of vertices $V_{0}, \ldots, V_{k}, k \geq 0$, that is a path in the underlying graph $H$ of $G ; k$ is called the length of the chain. A loop in $G$ is a sequence of vertices that is a cycle in the underlying graph $H$ of $G$.

Note that, in a digraph, the concept of path takes the directions of the arcs into account, while the concept of chain does not. A digraph is therefore acyclic if it contains no directed cycles; an acyclic digraph (or DAG) can contain loops, however.

In a directed graph, two vertices may be connected by a chain. If this property holds for any two vertices in a digraph, we say that the graph is connected.

Definition 2.1.6 $A$ digraph $G$ is connected if there exists at least one chain between any two vertices in $G$; otherwise, it is called unconnected.

We have introduced the concept of connectedness to apply to directed graphs; the concept, however, is easily extended to apply to undirected graphs.

We now distinguish between several types of digraph.
Definition 2.1.7 $A$ digraph $G$ is called singly connected if does not contain any loops; otherwise, it is called multiply connected. A singly connected digraph $G$ is called a directed tree if each vertex in $G$ has at most one predecessor.

Note that in a singly connected digraph, there is at most one chain between any two vertices; this property does not hold for multiply connected digraphs.

To conclude, we introduce the concept of a subgraph. The concept is introduced for undirected graphs, but is extended straightforwardly to apply to digraphs.

Definition 2.1.8 Let $G=\left(\boldsymbol{V}_{G}, \boldsymbol{E}_{G}\right)$ be an undirected graph. The subgraph $H$ induced by $\boldsymbol{V} \subseteq \boldsymbol{V}_{G}$ is the undirected graph $H=\left(\boldsymbol{V},(\boldsymbol{V} \times \boldsymbol{V}) \cap \boldsymbol{E}_{G}\right)$.

Note that a subgraph induced by a set of vertices $\boldsymbol{V}$ takes from the original graph all edges existent among the vertices from $\boldsymbol{V}$.

### 2.2 Probability Theory

In this section, we provide a brief review of some basic concepts from probability theory. Once more, our review is tailored to the Bayesian network framework and is not meant to be an exhaustive tutorial; for further information, any introductory textbook on probability theory will suffice.

Probability theory is often approached from a set-theoretic point of view. Probability distributions are then defined on sets of elements that represent events. All possible outcomes of an experiment (for example, the possible outcomes of rolling a die) are given by the sample space $\Omega$ and each event $A$ is a subset of $\Omega$. A probability measure/function/distribution then is a function from events to the $[0,1]$ interval. As events are sets, combinations of events can be expressed using operations on sets such as union ( $\cup$ ) and intersection ( $\cap$ ). Outcomes of an experiment are often coded by using a random variable (also: statistical/ stochastic variable) which is a function from the sample space to another space (such as reals). By writing probability distributions on statistical variables, the notation suppresses references to the actual sample space. However, as a statement about a statistical variable defines an event, there is no actual difference.

In the Bayesian network community random variables are taken to be functions from $\Omega$ to $\Omega$. For a variable $V$ defined on outcomes true and false, for example, we therefore have $V($ true $)=$ true and $V($ false $)=$ false. We now simply say that $V$ can have or take on one of the values true and false, in which case we write $V=$ true or $V=$ false as possible value-assignments. The Bayesian network community now approaches probability theory from an algebraic point of view by associating probabilities with logical propositions instead of with sets.

In this syllabus, we consider a set of statistical variables $\boldsymbol{V}=\left\{V_{1}, \ldots, V_{n}\right\}, n \geq 1$. For ease of exposition, we will often restrict the discussion to binary variables taking one of the truth values true and false; the generalisation to variables with more than two discrete values, however, is rather straightforward. For abbreviation, we will use $v_{i}$ to denote the proposition that the variable $V_{i}$ takes the value true; $V_{i}=$ false will be denoted as $\neg v_{i}$. The set of variables $\boldsymbol{V}$ may be looked upon as spanning a Boolean algebra of propositions $\mathcal{V}$. Informally speaking, this algebra comprises all logical propositions that are built from value assignments to the variables discerned. More formally, the Boolean algebra of propositions $\mathcal{V}$ spanned by $\boldsymbol{V}$ is the set of logical propositions consisting of the constant propositions True and False ${ }^{2}$, the atomic proposition $v_{i}$ for all $V_{i} \in \boldsymbol{V}$, and all compound propositions that are constructed from these by applying the binary operators $\wedge($ conjunction $), \vee($ disjunction $)$, and the unary operator $\neg$ (negation); the elements of the algebra $\mathcal{V}$ adhere to the usual axioms of propositional logic.

We now define a joint probability distribution as a function on a Boolean algebra of propositions that is spanned by a set of statistical variables.

[^1]Definition 2.2.1 Let $\boldsymbol{V}$ be a set of statistical variables and let $\mathcal{V}$ be the Boolean algebra of propositions spanned by $\boldsymbol{V}$. Let $\operatorname{Pr}: \mathcal{V} \rightarrow[0,1]$ be a function such that

- $\operatorname{Pr}(x) \geq 0$, for all $x \in \mathcal{V}$, and $\operatorname{Pr}($ False $)=0$, more in specific;
- $\operatorname{Pr}($ True $)=1$;
- $\operatorname{Pr}(x \vee y)=\operatorname{Pr}(x)+\operatorname{Pr}(y)$, for all $x, y \in \mathcal{V}$ such that $x \wedge y \equiv$ False.

Then, $\operatorname{Pr}$ is called a joint probability distribution on $\boldsymbol{V}$. For each $x \in \mathcal{V}$, the function value $\operatorname{Pr}(x)$ is termed the probability of $x$.

A probability $\operatorname{Pr}(x)$ for a logical proposition $x$ expresses the amount of certainty concerning the truth of $x$. Note that in the previous definition we have associated probabilities with logical propositions instead of with sets, which is the more common view taken in (introductory) literature on probability theory. It can easily be shown, however, that the probability of an event (a set of outcomes) is equivalent to the probability of the truth of the proposition asserting the occurrence of the event [Finetti, 1970].

Example 2.2.2 Suppose $X$ and $Y$ are statistical variables representing a coin toss. Let $A$ be the event that $X=$ heads and $Y=$ tails then the probability of this event is

- from a set-theoretic point of view: the probability of event $A$, written $\operatorname{Pr}(A)$;
- from an algebraic point of view: the probability that $X=$ heads and $Y=$ tails are both true propositions, written $\operatorname{Pr}(X=$ heads $\wedge Y=$ tails $)$

If $X=$ heads and $Y=$ tails were considered two separate events $A$ and $B$, then this would make no difference from the algebraic point of view, but in the set-theoretic approach we should now write $\operatorname{Pr}(A \cap B)$.

In the sequel, we will want to single out strictly positive joint probability distributions as these have some interesting properties. Strictly positive distributions for example are well-known for their not embedding any functional or logical relationships among their variables.

Definition 2.2.3 Let $\boldsymbol{V}$ be a set of statistical variables and let $\mathcal{V}$ be the Boolean algebra of propositions spanned by $\boldsymbol{V}$. Let $\operatorname{Pr}$ be a joint probability distribution on $\boldsymbol{V}$. $\operatorname{Pr}$ is strictly positive if $\operatorname{Pr}(x)=0$ implies $x \equiv$ False.

We now introduce the concept of conditional probability.

Definition 2.2.4 Let $\boldsymbol{V}$ be a set of statistical variables and let $\mathcal{V}$ be the Boolean algebra of propositions spanned by $\boldsymbol{V}$. Let $\operatorname{Pr}$ be a joint probability distribution on $\boldsymbol{V}$. For each $x, y \in \mathcal{V}$ with $\operatorname{Pr}(y)>0$, the conditional probability of $x$ given $y$, denoted as $\operatorname{Pr}(x \mid y)$, is defined as

$$
\operatorname{Pr}(x \mid y)=\frac{\operatorname{Pr}(x \wedge y)}{\operatorname{Pr}(y)}
$$

The conditional probability $\operatorname{Pr}(x \mid y)$ expresses the amount of certainty concerning the truth of $x$ given that the information $y$ is known with certainty. Note that a conditional probability $\operatorname{Pr}(x \mid y)=p$ does not mean that whenever $y$ is known to be true, the probability of $x$ equals $p$ : it means that the probability of $x$ equals $p$ if $y$ is known to be true and nothing else is known that may affect the certainty concerning the truth of $x$. In the sequel, we will assume that all conditional probabilities specified are properly defined, that is, for each conditional probability $\operatorname{Pr}(x \mid y)$, we will implicitly assume that $\operatorname{Pr}(y)>0$. We further state without proof that for a given element $y \in \mathcal{V}$, the conditional probabilities $\operatorname{Pr}(x \mid y)$ for all $x \in \mathcal{V}$ once more constitute a joint probability distribution on $\boldsymbol{V}$; this probability distribution is called the conditional probability distribution given $y$ and will sometimes be denoted as $\operatorname{Pr}^{y}$. A conditional probability distribution is sometimes referred to as a posterior probability distribution; the joint probability distribution it is obtained from then in contrast is referred to as the prior distribution.

The following definition introduces the concept of independence of propositions.
Definition 2.2.5 Let $\boldsymbol{V}$ be a set of statistical variables and let $\mathcal{V}$ be the Boolean algebra of propositions spanned by $\boldsymbol{V}$. Let $\operatorname{Pr}$ be a joint probability distribution on $\boldsymbol{V}$. Then, two propositions $x, y \in \mathcal{V}$ are called (mutually) independent in $\operatorname{Pr}$ if

$$
\operatorname{Pr}(x \wedge y)=\operatorname{Pr}(x) \cdot \operatorname{Pr}(y)
$$

otherwise, $x$ and $y$ are called dependent in $\operatorname{Pr}$. Two propositions $x, y \in \mathcal{V}$ are called conditionally independent given the proposition $z \in \mathcal{V}$ in $\operatorname{Pr}$ if

$$
\operatorname{Pr}(x \wedge y \mid z)=\operatorname{Pr}(x \mid z) \cdot \operatorname{Pr}(y \mid z)
$$

otherwise, $x$ and $y$ are called conditionally dependent given $z$ in $\operatorname{Pr}$.
In the sequel, we will make extensive use of various well-known properties of joint probability distributions. Before stating these properties, we provide some additional concepts and notational conventions. Recall that so far we have built on the Boolean algebra of propositions $\mathcal{V}$ spanned by some set of statistical variables $\boldsymbol{V}$. In the sequel we will consider propositions from the Boolean algebra spanned by $\boldsymbol{V}$ without explicitly refering to $\mathcal{V}$. More specifically, we will typically consider (arbitrary) conjunctions of value assignments to all variables from the set $\boldsymbol{V}$ or from some subset $\boldsymbol{W}$ of $\boldsymbol{V}$. Rather than describing a joint probability distribution on $\boldsymbol{V}$ in terms of $\operatorname{Pr}(x)$ for any proposition $x \in \mathcal{V}$, we write $\operatorname{Pr}(\boldsymbol{V})$ to denote the collection of probabilities $\operatorname{Pr}\left(c_{\boldsymbol{V}}\right)$ where $c_{V}$ is a conjunction of value assignments to the variables from $\boldsymbol{V}$, called a configuration. Statements involving a set of variables $\boldsymbol{V}$ are thus taken to hold for any configuration $c_{V}$, where a configuration $c_{V}$ is nothing more than a short-hand notation to indicate that you are considering a proposition that consists of the conjunction of atomic propositions representing some value assignment to each variable in the set $\boldsymbol{V} .{ }^{3}$ Sometimes we will write $\boldsymbol{v}$ to indicate one specific configuration of $\boldsymbol{V}$. If a configuration $c_{\boldsymbol{W}}$ concerns the empty set $\boldsymbol{W}=\varnothing$, then-by convention- $c_{\boldsymbol{W}} \equiv$ True (no restrictions, anything is possible). To avoid abundance of braces, we will often write $c_{V_{i}}$ instead of $c_{\left\{V_{i}\right\}}$ for singleton sets $\left\{V_{i}\right\}$.

[^2]We now state the various properties that we will use in the sequel. We would like to note that in the literature on probability theory these properties are introduced in many different appearances; we have chosen the form that suits our purposes best. The property stated in the following proposition is known as the chain rule.

Proposition 2.2.6 Let $\boldsymbol{V}=\left\{V_{1}, \ldots, V_{n}\right\}, n \geq 1$, be a set of statistical variables and let $\operatorname{Pr}$ be a joint probability distribution on $\boldsymbol{V}$. Then,

$$
\operatorname{Pr}(\boldsymbol{V})=\operatorname{Pr}\left(V_{1} \wedge \cdots \wedge V_{n}\right)=\operatorname{Pr}\left(V_{n} \mid V_{1} \wedge \cdots \wedge V_{n-1}\right) \cdot \ldots \cdot \operatorname{Pr}\left(V_{2} \mid V_{1}\right) \cdot \operatorname{Pr}\left(V_{1}\right)
$$

Note that the expression stated in the previous proposition is actually a template for a whole collection of equalities: the conjunctions only form proper propositions once the variables are assigned a value ( $V_{1} \wedge V_{2}$ is not a proposition since $V_{1}$ and $V_{2}$ are not propositions, but $V_{1}=\operatorname{red} \wedge V_{2}=$ yellow is). For example, if all variables are binaryvalued, taking on either the value true or the value false, the expression represents $2^{n}$ equalities, one for each configuration of the set of variables $\boldsymbol{V}$.

The property stated in the following proposition is termed the marginalisation property.

Proposition 2.2.7 Let $\boldsymbol{V}$ be a set of statistical variables and let $\operatorname{Pr}$ be a joint probability distribution on $\boldsymbol{V}$. Then,

$$
\operatorname{Pr}(\boldsymbol{X})=\sum_{c_{\boldsymbol{Y}}} \operatorname{Pr}\left(\boldsymbol{X} \wedge c_{\boldsymbol{Y}}\right)
$$

for all sets of variables $\boldsymbol{X}, \boldsymbol{Y} \subseteq \boldsymbol{V}$.
We state without proof that for any set of variables $\boldsymbol{X} \subseteq \boldsymbol{V}$, the probabilities $\operatorname{Pr}\left(c_{\boldsymbol{X}}\right)$ for all configurations $c_{\boldsymbol{X}}$ of $\boldsymbol{X}$ once more constitute a joint probability distribution; this probability distribution is termed the marginal probability distribution on $\boldsymbol{X}$.

The conditioning property is stated in the following proposition.
Proposition 2.2.8 Let $\boldsymbol{V}$ be a set of statistical variables and let $\operatorname{Pr}$ be a joint probability distribution on $\boldsymbol{V}$. Then,

$$
\operatorname{Pr}(\boldsymbol{X})=\sum_{c_{\boldsymbol{Y}}} \operatorname{Pr}\left(\boldsymbol{X} \mid c_{\boldsymbol{Y}}\right) \cdot \operatorname{Pr}\left(c_{\boldsymbol{Y}}\right)
$$

for all sets of variables $\boldsymbol{X}, \boldsymbol{Y} \subseteq \boldsymbol{V}$.
The following theorem is known as Bayes' Theorem.
Theorem 2.2.9 Let $\boldsymbol{V}$ be a set of statistical variables and let $\operatorname{Pr}$ be a joint probability distribution on $\boldsymbol{V}$. Then,

$$
\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Y})=\frac{\operatorname{Pr}(\boldsymbol{Y} \mid \boldsymbol{X}) \cdot \operatorname{Pr}(\boldsymbol{X})}{\operatorname{Pr}(\boldsymbol{Y})}
$$

for all sets of variables $\boldsymbol{X}, \boldsymbol{Y} \subseteq \boldsymbol{V}$.
To conclude this section, we once more turn to the concept of (conditional) independence. Recall that so far we have taken the concept of independence to apply to propositions. We now introduce the concept of independence of variables.

Definition 2.2.10 Let $\boldsymbol{V}$ be a set of statistical variables and let $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}$. Let $\operatorname{Pr}$ be a joint probability distribution on $\boldsymbol{V}$. Then, the set of variables $\boldsymbol{X}$ is called conditionally independent of the set of variables $\boldsymbol{Y}$ given the set of variables $\boldsymbol{Z}$ in $\operatorname{Pr}$ if

$$
\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Y} \wedge \boldsymbol{Z})=\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z})
$$

otherwise, $\boldsymbol{X}$ is called conditionally dependent of $\boldsymbol{Y}$ given $\boldsymbol{Z}$ in $\operatorname{Pr}$.

In qualitative terms, the expression $\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Y} \wedge \boldsymbol{Z})=\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z})$ indicates that, once information about $\boldsymbol{Z}$ is available, information about $\boldsymbol{Y}$ is irrelevant with respect to $\boldsymbol{X}$. Note that for $\boldsymbol{X}$ and $\boldsymbol{Y}$ to be independent given $\boldsymbol{Z}$, every pair of configurations of $\boldsymbol{X}$ and $\boldsymbol{Y}$ has to be independent given every configuration of $\boldsymbol{Z}$. Independence of variables therefore implies independence of propositions. The reverse property, however, does not hold in general. Also note that the expression from Definition 2.2.10 is asymmetric in $\boldsymbol{X}$ and $\boldsymbol{Y}$. Using Bayes' Theorem, however, it is easily shown that $\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Y} \wedge \boldsymbol{Z})=$ $\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z})$ implies $\operatorname{Pr}(\boldsymbol{Y} \mid \boldsymbol{X} \wedge \boldsymbol{Z})=\operatorname{Pr}(\boldsymbol{Y} \mid \boldsymbol{Z})$.

### 2.3 Notation

In this section we summarise our notation. The attentive reader will have already spotted most of the conventions used both implicitly and explicitly in the above two sections.

We use upper case letters to denote a single variable or vertex $V$, and lower case to indicate a specific configuration $v$ of variable $V$. Arbitrary configurations are denoted $c_{V}$; for binary-valued variables, taking on either the value true or the value false, the two configurations are typically denoted $v_{i}$ and $\neg v_{i}$, respectively. Boldfaced capital letters $\boldsymbol{V}$ are used to indicate sets of vertices or variables. An arbitrary configuration for $\boldsymbol{V}$ is denoted $c_{\boldsymbol{V}}$, whereas a boldface lower case $\boldsymbol{v}$ indicates a specific configuration for a set of variables $\boldsymbol{V}$.

Upper case letters are also used to represent templates of all possible configurations; the distinction between a set and a template should be clear from the context.

Example 2.3.1 Consider the set of two binary-valued statistical variables $\boldsymbol{V}=\{X, Y\}$. The joint probability distribution on $\boldsymbol{V}$ is written as $\operatorname{Pr}(\boldsymbol{V})$, or $\operatorname{Pr}(X \wedge Y)$ where $\boldsymbol{V}$ and $X \wedge Y$ are actually templates that represent all 4 possible configurations $c_{\boldsymbol{V}} \equiv c_{X} \wedge c_{Y}$ : $x \wedge y, x \wedge \neg y, \neg x \wedge y$ and $\neg x \wedge \neg y$. The joint distributions are completely defined by the probabilities $\operatorname{Pr}\left(c_{\boldsymbol{V}}\right)=\operatorname{Pr}\left(c_{X} \wedge c_{Y}\right)$. Let $v \equiv x \wedge \neg y$ be a specific configuration, then $\operatorname{Pr}(v)=\operatorname{Pr}(x \wedge \neg y)$.

The conditional distribution on $X$ given $Y=$ false is written as $\operatorname{Pr}(X \mid \neg y)$ where again $X$ is a template representing the two configurations $c_{X} \cdot \operatorname{Pr}(X \mid Y)$ indicates a set of conditional distributions on $X$ : one for each configuration $c_{Y}$ of $Y$.

## Exercises

## Exercise 2.1

Prove the following properties for any joint probability distribution, using only definitions and not the properties from this exercise:
a. the chain rule (stated in Proposition 2.2.6);
b. Bayes' Theorem (stated in Theorem 2.2.9);
c. the marginalisation property (stated in Proposition 2.2.7);

* d. the conditioning property (stated in Proposition 2.2.8).


## Exercise 2.2

Let $\boldsymbol{V}$ be a set of statistical variables and let $\operatorname{Pr}$ be a joint probability distribution on $\boldsymbol{V}$. Show that

$$
\operatorname{Pr}(\boldsymbol{X} \vee \boldsymbol{Y})=\operatorname{Pr}(\boldsymbol{X})+\operatorname{Pr}(\boldsymbol{Y})-\operatorname{Pr}(\boldsymbol{X} \wedge \boldsymbol{Y})
$$

for all sets of variables $\boldsymbol{X}, \boldsymbol{Y} \subseteq \boldsymbol{V}$.

## Exercise 2.3

Let $\boldsymbol{V}$ be a set of statistical variables and let $\operatorname{Pr}$ be a joint probability distribution on $\boldsymbol{V}$. Show that

$$
\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z})=\sum_{c_{\boldsymbol{Y}}} \operatorname{Pr}\left(\boldsymbol{X} \mid c_{\boldsymbol{Y}} \wedge \boldsymbol{Z}\right) \cdot \operatorname{Pr}\left(c_{\boldsymbol{Y}} \mid \boldsymbol{Z}\right)
$$

for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}$.

* Exercise 2.4

Let $\boldsymbol{V}$ be a set of statistical variables and let $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}$. Let $\operatorname{Pr}$ be a joint probability distribution on $\boldsymbol{V}$. Show that the set of variables $\boldsymbol{X}$ is conditionally independent of the set of variables $\boldsymbol{Y}$ given the set of variables $\boldsymbol{Z}$ if and only if $\operatorname{Pr}(\boldsymbol{X} \wedge \boldsymbol{Y} \mid \boldsymbol{Z})=\operatorname{Pr}(\boldsymbol{X} \mid$ $\boldsymbol{Z}) \cdot \operatorname{Pr}(\boldsymbol{Y} \mid \boldsymbol{Z})$.

## Chapter 3

## Independences and Graphical Representations


#### Abstract

The historical background to the framework of Bayesian networks shows various attempts to handle the computational complexity of applying probability theory for reasoning with uncertainty in knowledge-based systems. The concept of (conditional) independence plays a key role in these attempts as knowledge about independences allows for simplifying computations. In this chapter, we address formalisms that allow for a concise representation of an independence relation for effective use in a knowledge-based system. Current research into independence relations (beyond the scope of this syllabus) is focussed on defining small generating sets [Waal \& Van der Gaag, 2005], and on automated construction of graphical representations from them [Baioletti et al., 2011].


### 3.1 The Concept of Independence Revisited

In most introductory literature on probability theory, the concept of (conditional) independence is introduced in terms of numerical quantities: the independence relation of a joint probability distribution is taken to be implicitly embedded in the probabilities involved. Recall for example that in the previous chapter we have defined two sets of variables $\boldsymbol{X}$ and $\boldsymbol{Y}$ to be conditionally independent given a third set of variables $\boldsymbol{Z}$ if $\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Y} \wedge \boldsymbol{Z})=\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z})$. A definition of independence in terms of numbers suggests that, in order to determine whether two sets of variables are (conditionally) independent, several conditional probabilities have to be computed and several equalities have to be tested; moreover, such a definition suggests that for determining independence a joint probability distribution has to be explicitly available for the variables discerned. In contrast, humans tend to be able to state directly, with conviction and consistency, whether or not two sets of variables are independent. Such statements of independence typically are issued qualitatively, without any reference to numerical manipulation of exact probabilities. Based on these observations, we cannot but conclude that the concept of independence is far more basic to human reasoning than its numerical definition suggests. In fact, the definition of independence in terms of probabilities may be looked upon as a quantitative way of capturing the basic concept which is qualitative in nature. To formalise properties of the qualitative concept of independence, J. Pearl and his co-researchers have designed an axiomatic system for independence [Pearl \& Paz, 1985, Pearl \& Verma, 1987, Geiger \& Pearl, 1988]. In this
section, we review this axiomatic system.

### 3.1.1 Pearl's Axiomatic System for Independence

We begin our review of Pearl's axiomatic system for independence by introducing some new terminology and notational convention.

Definition 3.1.1 Let $\boldsymbol{V}$ be a set of statistical variables and let $\operatorname{Pr}$ be a joint probability distribution on $\boldsymbol{V}$. Then, the independence relation $I_{\operatorname{Pr}} \subseteq \mathcal{P}(\boldsymbol{V}) \times \mathcal{P}(\boldsymbol{V}) \times \mathcal{P}(\boldsymbol{V})$ of Pr is defined by $(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \in I_{\operatorname{Pr}}$ if and only if $\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Y} \wedge \boldsymbol{Z})=\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z})$ for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}$.

In the sequel, we will write $I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y})$ to denote $(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \in I_{\operatorname{Pr}}$ and $\neg I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y})$ to denote $(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \notin I_{\operatorname{Pr}}$. A statement $I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y})$ of a joint probability distribution's independence relation $I_{\mathrm{Pr}}$ is termed an independence statement. In qualitative terms, an independence statement $I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y})$ expresses that in the context of information about $\boldsymbol{Z}$ information about $\boldsymbol{Y}$ is irrelevant with respect to $\boldsymbol{X}$. We note that the above definition allows for stating some trivial but convenient independence statements, such as $I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{X}, \boldsymbol{Y})$, which holds iff $\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Y} \wedge \boldsymbol{X})=\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{X})$, i.e. $1=1$; its symmetric version $I_{\operatorname{Pr}}(\boldsymbol{Y}, \boldsymbol{X}, \boldsymbol{X})$ is also trivially true since $I_{\operatorname{Pr}}(\boldsymbol{Y}, \boldsymbol{X}, \boldsymbol{X})$ iff $\operatorname{Pr}(\boldsymbol{Y} \mid \boldsymbol{X} \wedge \boldsymbol{X})=\operatorname{Pr}(\boldsymbol{Y} \mid \boldsymbol{X})$.

In designing his axiomatic system for independence, Pearl builds on a set of properties that are satisfied by any joint probability distribution's independence relation; Theorem 3.1.2 reviews these properties.

Theorem 3.1.2 Let $\boldsymbol{V}$ be a set of statistical variables. Let $\operatorname{Pr}$ be a joint probability distribution on $\boldsymbol{V}$ and let $I_{\operatorname{Pr}}$ be its independence relation. Then, $I_{\operatorname{Pr}}$ satisfies the properties

- $I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \rightarrow I_{\mathrm{Pr}}(\boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{X})$;
(symmetry)
- $I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \rightarrow I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \wedge I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{W}) ; \quad$ (decomposition)
- $I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \rightarrow I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{W}, \boldsymbol{Y}) ; \quad$ (weak union)
- $I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \wedge I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{Y}, \boldsymbol{W}) \rightarrow I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) ; \quad$ (contraction)
for all mutually disjoint sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W} \subseteq \boldsymbol{V}$. If the distribution $\operatorname{Pr}$ is strictly positive, then $I_{\operatorname{Pr}}$ satisfies the additional property

$$
\text { - } I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{W}, \boldsymbol{Y}) \wedge I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{Y}, \boldsymbol{W}) \rightarrow I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) ; \quad \text { (intersection) }
$$

for all mutually disjoint sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W} \subseteq \boldsymbol{V}$.
The properties stated in the previous theorem are easily verified from the basic axioms of probability theory. We would like to note that we have closely followed Pearl by stating the properties in the theorem to hold for mutually disjoint sets of variables only [Pearl, 1988]. These properties, however, also hold for overlapping sets of variables [Van der Gaag \& Meyer, 1998].

Pearl now takes the properties stated in Theorem 3.1.2 as axioms for the qualitative concept of independence [Pearl, 1988]. Following the properties for $I_{\mathrm{Pr}}$, Pearl assumed for each axiom that the sets of variables involved are mutually disjoint. Given the
insight that the properties also hold for overlapping sets, we will lift the assumption of mutual disjointness in the next and all following definitions involving independence relations. The following now defines informational independence:

Definition 3.1.3 Let $\boldsymbol{V}$ be a set of statistical variables. A semi-graphoid independence relation on $\boldsymbol{V}$ is a ternary relation $I \subseteq \mathcal{P}(\boldsymbol{V}) \times \mathcal{P}(\boldsymbol{V}) \times \mathcal{P}(\boldsymbol{V})$ such that $I$ satisfies the properties

- $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \rightarrow I(\boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{X})$;
- $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \rightarrow I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \wedge I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{W})$;
- $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \rightarrow I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{W}, \boldsymbol{Y})$;
- $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \wedge I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{Y}, \boldsymbol{W}) \rightarrow I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W})$;
for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W} \subseteq \boldsymbol{V}$. A graphoid independence relation $I$ on $\boldsymbol{V}$ is a semi-graphoid independence relation on $\boldsymbol{V}$ such that I satisfies the additional property
- $I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{W}, \boldsymbol{Y}) \wedge I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{Y}, \boldsymbol{W}) \rightarrow I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) ;$
for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W} \subseteq \boldsymbol{V}$.
The properties described in the previous definition with each other convey the idea that learning irrelevant information does not alter the independences among the variables discerned [Pearl, 1988]. We consider the qualitative meanings of the various properties separately.

The property

$$
I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \rightarrow I(\boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{X})
$$

for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}$, states that if information about $\boldsymbol{Y}$ is deemed irrelevant with respect to $\boldsymbol{X}$ in the context of some information about $\boldsymbol{Z}$, then information about $\boldsymbol{X}$ must be irrelevant with respect to $\boldsymbol{Y}$ in this context; this property is called the symmetry axiom.

The property

$$
I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \rightarrow I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \wedge I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{W})
$$

for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W} \subseteq \boldsymbol{V}$, asserts that if information about both $\boldsymbol{Y}$ and $\boldsymbol{W}$ is judged irrelevant with respect to $\boldsymbol{X}$, then both information about $\boldsymbol{Y}$ and information about $\boldsymbol{W}$ must be irrelevant with respect to $\boldsymbol{X}$ separately; this property is known as the decomposition axiom. We would like to note that the decomposition axiom may be reformulated as $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \rightarrow I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y})$; we have chosen, however, to use Pearl's original formulation because it conveys the idea of decomposition more clearly.

The property

$$
I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \rightarrow I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{W}, \boldsymbol{Y})
$$

for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W} \subseteq \boldsymbol{V}$, states that learning information about $\boldsymbol{W}$ that is known to be irrelevant with respect to $\boldsymbol{X}$ cannot help irrelevant information
about $\boldsymbol{Y}$ to become relevant with respect to $\boldsymbol{X}$; this property is known as the weak union axiom.

The property

$$
I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \wedge I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{Y}, \boldsymbol{W}) \rightarrow I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W})
$$

for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W} \subseteq \boldsymbol{V}$, states that if we judge information about $\boldsymbol{W}$ to be irrelevant with respect to $\boldsymbol{X}$ after learning some irrelevant information about $\boldsymbol{Y}$, then the information about $\boldsymbol{W}$ must have been irrelevant with respect to $\boldsymbol{X}$ before we learned $\boldsymbol{Y}$; this property is known as the contraction axiom. Note that the contraction axiom can be reformulated as $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \rightarrow(I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{Y}, \boldsymbol{W}) \rightarrow I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}))$. From this reformulation, it is seen that the axiom can be looked upon as a conditional reverse of the weak union axiom.

We now consider the property

$$
I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{W}, \boldsymbol{Y}) \wedge I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{Y}, \boldsymbol{W}) \rightarrow I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W})
$$

for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W} \subseteq \boldsymbol{V}$, for graphoid independence relations. This property states that if, in the context of some information about $\boldsymbol{Z}$, learning information about $\boldsymbol{W}$ renders information about $\boldsymbol{Y}$ irrelevant with respect to $\boldsymbol{X}$ and learning $\boldsymbol{Y}$ renders $\boldsymbol{W}$ irrelevant with respect to $\boldsymbol{X}$, then the information about both $\boldsymbol{Y}$ and $\boldsymbol{W}$ must be irrelevant with respect to $\boldsymbol{X}$ given $\boldsymbol{Z}$; this property is known as the intersection axiom.

From Definition 3.1.3 and Theorem 3.1.2, we observe that every independence relation that is embedded in a joint probability distribution is a semi-graphoid independence relation; this property is stated more formally in the following corollary.

Corollary 3.1.4 Let $\boldsymbol{V}$ be a set of statistical variables. Let $\operatorname{Pr}$ be a joint probability distribution on $\boldsymbol{V}$ and let $I_{\operatorname{Pr}}$ be its independence relation. Then, $I_{\operatorname{Pr}}$ is a semi-graphoid independence relation. Furthermore, if $\operatorname{Pr}$ is strictly positive, then $I_{\operatorname{Pr}}$ is a graphoid independence relation.

Unfortunately, although any probability distribution's independence relation is a semigraphoid independence relation, the reverse property does not hold. There exist semigraphoid independence relations for which there do not exist joint probability distributions embedding them; for details, we refer to [Van der Gaag \& Meyer, 1996, Studený, 1989]. We would like to note that it has been shown that a finite axiomatisation of the concept of probabilistic independence does not exist [Studený, 1992].

### 3.1.2 Properties of Independence Relations

Using the definition of informational independence, we derive some convenient properties of (semi-graphoid and graphoid) independence relations. The following lemma shows that the symmetry and contraction axioms are easily generalised to bi-implications.

Lemma 3.1.5 Let $\boldsymbol{V}$ be a set of statistical variables. Furthermore, let I be a semigraphoid independence relation on $\boldsymbol{V}$. Then,

- $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \leftrightarrow I(\boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{X}) ;$
- $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \wedge I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{Y}, \boldsymbol{W}) \leftrightarrow I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) ;$
for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W} \subseteq \boldsymbol{V}$.
Proof. We begin our proof by observing that since $I$ is a semi-graphoid independence relation, it obeys the first four axioms stated in Definition 3.1.3. The first property stated in the lemma now follows directly from the symmetry axiom. For the second property, we observe that $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \wedge I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{Y}, \boldsymbol{W}) \rightarrow I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W})$ coincides with the contraction axiom and therefore trivially holds for the relation $I$. We will now prove that $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \rightarrow I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \wedge I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{Y}, \boldsymbol{W})$. We have

$$
I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \Rightarrow I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \wedge I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{W}) \Rightarrow I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y})
$$

by the decomposition axiom. In addition, we have

$$
I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \Rightarrow I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{Y}, \boldsymbol{W})
$$

by weak union. The property stated in the lemma now follows directly.
For graphoid independence relations we have that the intersection axiom can also be generalised to a bi-implication.

Lemma 3.1.6 Let $\boldsymbol{V}$ be a set of statistical variables. Furthermore, let I be a graphoid independence relation on $\boldsymbol{V}$. Then,

$$
I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{W}, \boldsymbol{Y}) \wedge I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{Y}, \boldsymbol{W}) \leftrightarrow I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W})
$$

for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W} \subseteq \boldsymbol{V}$.
Proof. We will only prove that $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \rightarrow I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{W}, \boldsymbol{Y}) \wedge I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{Y}, \boldsymbol{W})$; the reverse property coincides with the intersection axiom and therefore trivially holds for the independence relation $I$. We have

$$
I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \Rightarrow I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{W}, \boldsymbol{Y})
$$

and

$$
I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \Rightarrow I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{Y}, \boldsymbol{W})
$$

by the weak union axiom. The property stated in the lemma now follows directly.
In the sequel, we will use the phrase independence relation to denote a semi-graphoid independence relation, unless stated otherwise.

### 3.2 Graphical Representations of Independence

One of the problems in applying probability theory for automated reasoning with uncertainty in a knowledge-based system is the space complexity of representing a joint probability distribution. Since the concept of independence plays a key role in solving this problem, a formalism for representing joint probability distributions should allow for efficiently modeling independences. There are various ways of representing an independence relation. One way, for example, is to enumerate the separate statements of
an independence relation explicitly. Such a representation clearly is impractical as the number of tuples in an independence relation can be astronomical. Another way is to make use of the axioms from Definition 3.1.3: only the statements from an appropriate subset of the independence relation are enumerated explicitly and all its other statements are defined implicitly by this set and the defining axioms. Although exploiting the axioms allows for a far more economical representation of an independence relation than explicit enumeration, it can still require exponential space. In this section, we consider more concise representations of independence relations, building on the idea of graphical encoding. In Section 3.2.1 we address modeling an independence relation in an undirected graph; in Section 3.2.2 we consider the representation of an independence relation in the formalism of directed graphs.

### 3.2.1 Undirected Graphs

Undirected graphs have no probabilistic meaning by themselves. For representing an independence relation in an undirected graph, therefore, a probabilistic meaning has to be assigned to the topological properties of such a graph, that is, we have to assign a probabilistic meaning to the vertices of the graph and to assign a probabilistic meaning to its edges. Informally speaking, we choose to encode an independence relation in an undirected graph by modelling the variables of the relation as vertices and by representing its independence statements by absence of edges. To formally capture this meaning, we begin by defining a graphical criterion for reading from a graph sets of vertices that allow for blocking all paths between two given sets of vertices; this graphical criterion is termed the separation criterion for undirected graphs.

Definition 3.2.1 Let $G=\left(\boldsymbol{V}_{G}, \boldsymbol{E}_{G}\right)$ be an undirected graph. Let $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}_{G}$ be sets of vertices in $G$. The set of vertices $\boldsymbol{Z}$ is said to separate the sets of vertices $\boldsymbol{X}$ and $\boldsymbol{Y}$ in $G$, denoted as $\langle\boldsymbol{X}| \boldsymbol{Z}|\boldsymbol{Y}\rangle_{G}$, if for each vertex $V_{i} \in \boldsymbol{X}$ and each vertex $V_{j} \in \boldsymbol{Y}$ every simple path from $V_{i}$ to $V_{j}$ in $G$ contains at least one vertex from $\boldsymbol{Z}$.

We look upon a separating set of variables as effectively blocking influence: if a set of variables $\boldsymbol{Z}$ separates two sets of variables $\boldsymbol{X}$ and $\boldsymbol{Y}$, then $\boldsymbol{Z}$ is looked upon as blocking any flow of information or influence between $\boldsymbol{X}$ and $\boldsymbol{Y}$. Two sets of variables $\boldsymbol{X}$ and $\boldsymbol{Y}$ that are thus separated by a set $\boldsymbol{Z}$ now are taken to be conditionally independent given $\boldsymbol{Z}$.

Definition 3.2.2 Let $\boldsymbol{V}$ be a set of statistical variables and let $I$ be an independence relation on $\boldsymbol{V}$. Furthermore, let $G=\left(\boldsymbol{V}_{G}, \boldsymbol{E}_{G}\right)$ be an undirected graph with $\boldsymbol{V}_{G}=\boldsymbol{V}$. Then,

- the graph $G$ is called an undirected dependence map, or D-map for short, for $I$, if for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}$, we have: if $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y})$ then $\langle\boldsymbol{X}| \boldsymbol{Z}|\boldsymbol{Y}\rangle_{G}$;
- the graph $G$ is called an undirected independence map, or I-map for short, for $I$, if for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}$, we have: if $\langle\boldsymbol{X}| \boldsymbol{Z}|\boldsymbol{Y}\rangle_{G}$ then $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y})$;
- the graph $G$ is called an undirected perfect map, or P-map for short, for $I$, if $G$ is both an undirected D-map and an undirected I-map for $I$.


Figure 3.1: An Example Undirected D-map.

From the previous definition we have that in an undirected D-map for an independence relation any pair of neighbouring vertices represents a pair of variables that are dependent; however, not every pair of dependent variables needs be represented as a pair of neighbouring vertices. In an undirected D-map, therefore, a pair of non-neighbouring vertices may represent either a pair of dependent variables or a pair of independent variables.

Example 3.2.3 Let $\boldsymbol{V}=\left\{V_{1}, V_{2}, V_{3}, V_{4}\right\}$ be a set of statistical variables. We consider the independence relation $I$ on $\boldsymbol{V}$ that is defined by the independence statements $I\left(\left\{V_{1}\right\},\left\{V_{2}, V_{3}\right\},\left\{V_{4}\right\}\right)$ and $I\left(\left\{V_{2}\right\},\left\{V_{1}, V_{4}\right\},\left\{V_{3}\right\}\right)$. The undirected graph $G$ shown in Figure 3.1 is an example of an undirected D-map for $I$, since for each independence statement $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}), \boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}$, from the relation $I$, there exists a matching separation statement $\langle\boldsymbol{X}| \boldsymbol{Z}|\boldsymbol{Y}\rangle_{G}$. For example, for the independence statement $I\left(\left\{V_{1}\right\},\left\{V_{2}, V_{3}\right\},\left\{V_{4}\right\}\right)$ we find that the only path from $V_{1}$ to $V_{4}$ comprises the vertex $V_{2}$ which is included in the set $\left\{V_{2}, V_{3}\right\}$, that is, $\left\langle\left\{V_{1}\right\}\right|\left\{V_{2}, V_{3}\right\}\left|\left\{V_{4}\right\}\right\rangle_{G}$.

In an undirected I-map, any pair of non-neighbouring vertices represents a pair of variables that are independent; however, not every pair of independent variables is represented as a pair of non-neighbouring vertices. In an undirected I-map, therefore, a pair of neighbouring vertices may represent either a pair of dependent variables or a pair of independent variables.

Example 3.2.4 Consider once more the independence relation $I$ from Example 3.2.3. The graph $G$ shown in Figure 3.2 is an example of an undirected I-map for the relation $I$, since for each separation statement $\langle\boldsymbol{X}| \boldsymbol{Z}|\boldsymbol{Y}\rangle_{G}, \boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}$, read from $G$, there exists a matching independence statement $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y})$ : for the single separation statement $\left\langle\left\{V_{2}\right\}\right|\left\{V_{1}, V_{4}\right\}\left|\left\{V_{3}\right\}\right\rangle_{G}$ read from the graph, we find the statement $I\left(\left\{V_{2}\right\},\left\{V_{1}, V_{4}\right\},\left\{V_{3}\right\}\right)$ in the relation $I$.


Figure 3.2: An Example Undirected I-map.


Figure 3.3: An Example Undirected P-map.

An undirected P-map for an independence relation faithfully represents all independences as well as all dependences from the relation.

Example 3.2.5 Consider once more the independence relation $I$ from Example 3.2.3. The graph $G$ shown in Figure 3.3 is the only undirected P-map for $G$. Note that $G$ is an undirected D-map for $I$ as well as an undirected I-map for $I$.

From the above observations, we conclude that the various types of map provide for reading different types of probabilistic information from an undirected graph. If the graph is an undirected I-map for an independence relation, we can read independence statements from it as any separation statement is guaranteed to correspond with an independence statement; from an undirected D-map we can read dependence statements, and from an undirected P-map we can read both independence and dependence statements.

We now investigate the expressive power of the formalism of undirected graphs for representing independence relations.

Lemma 3.2.6 For every independence relation, there exist an undirected D-map and an undirected I-map.

Proof. We first show that for every independence relation there exists an undirected D-map. To this end, we show that the edgeless undirected graph $G=(\boldsymbol{V}, \varnothing)$ is an undirected D-map for any independence relation $I$ on the set of variables $\boldsymbol{V}$. Since the graph $G$ is edgeless, we have that the property: if $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y})$ then $\langle\boldsymbol{X}| \boldsymbol{Z}|\boldsymbol{Y}\rangle_{G}$, trivially holds for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}$. Informally speaking, as the graph $G$ does not represent any dependences, it cannot represent any dependence incorrectly.

We now show that for every independence relation there exists an undirected Imap. To this end, we show that the complete undirected graph $G^{\prime}=(\boldsymbol{V}, \boldsymbol{V} \times \boldsymbol{V})$ is an undirected I-map for any independence relation $I$ on $\boldsymbol{V}$. Since in this graph every pair of vertices is connected by an edge, we have that the property: if $\langle\boldsymbol{X}| \boldsymbol{Z}|\boldsymbol{Y}\rangle_{G}$ then $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y})$, trivially holds for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}$. Informally speaking, as the graph $G^{\prime}$ does not represent any independences, it cannot represent any independence incorrectly.

From the previous lemma, we have that every independence relation has an undirected D-map and an undirected I-map; in fact, an independence relation may have several undirected D-maps and I-maps. Unfortunately, a similar property does not hold for undirected P-maps: not every independence relation has an undirected P-map.

Example 3.2.7 Consider an independence relation $I$ on a set of variables $\boldsymbol{V}$ such that $I\left(\boldsymbol{X}, \boldsymbol{Z}_{\mathbf{1}}, \boldsymbol{Y}\right)$ for some sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}_{\mathbf{1}} \subseteq \boldsymbol{V}$ and $\neg I\left(\boldsymbol{X}, \boldsymbol{Z}_{\mathbf{1}} \cup \boldsymbol{Z}_{\mathbf{2}}, \boldsymbol{Y}\right)$ for some set $\boldsymbol{Z}_{\mathbf{2}} \subseteq \boldsymbol{V}$ with $\boldsymbol{Z}_{\mathbf{2}} \neq \varnothing$ and $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}_{\mathbf{1}}, \boldsymbol{Z}_{\mathbf{2}}$ are disjoint. The independence relation $I$ is an example of an independence relation comprising an induced dependence: the sets of variables $\boldsymbol{X}$ and $\boldsymbol{Y}$ are independent given $\boldsymbol{Z}_{\mathbf{1}}$, but become dependent if information about $\boldsymbol{Z}_{2}$ becomes available. Now, for any undirected graph $G$ to be an undirected P-map for this independence relation $I$, it should be an undirected D-map as well as an undirected I-map for $I$. For $G$ to be an undirected D-map for $I$, it should display the set $\boldsymbol{Z}_{1}$ as a separating set for $\boldsymbol{X}$ and $\boldsymbol{Y}$; on the other hand, for $G$ to be an undirected I-map, it should display $\boldsymbol{Z}_{\mathbf{1}} \cup \boldsymbol{Z}_{\mathbf{2}}$ as not separating $\boldsymbol{X}$ and $\boldsymbol{Y}$. We observe, however, that for any undirected graph $G$ the property: if $\langle\boldsymbol{X}| \boldsymbol{Z}_{\mathbf{1}}|\boldsymbol{Y}\rangle_{G}$ then $\langle\boldsymbol{X}| \boldsymbol{Z}_{\mathbf{1}} \cup \boldsymbol{Z}_{\mathbf{2}}|\boldsymbol{Y}\rangle_{G}$ holds by definition. More in general, if two sets of vertices $\boldsymbol{X}$ and $\boldsymbol{Y}$ are separated by a third set of vertices $\boldsymbol{Z}$ in an undirected graph $G$, then $\boldsymbol{X}$ and $\boldsymbol{Y}$ are separated by any superset of $\boldsymbol{Z}$ in $G$ as well. From this observation, it will be evident that no undirected graph can satisfy the two requirements mentioned above simultaneously and, hence, that there does not exist an undirected P-map for $I$.

Although not every independence relation has an undirected P-map, there are independence relations that can indeed be represented faithfully. An independence relation for which an undirected P-map exists is termed undirected graph-isomorphic.

Definition 3.2.8 An independence relation I is said to be undirected graph-isomorphic if there exists an undirected graph $G$ such that $G$ is an undirected $P$-map for $I$.

We would like to note that, if an independence relation is undirected graph-isomorphic, then it allows one and only one undirected P-map. To conclude, we would like to mention that a necessary and sufficient condition has been identified for an independence relation to be undirected graph-isomorphic [Pearl, 1988]; this condition allows for testing whether a given independence relation lends itself to representation in an undirected graph. Here, we will not elaborate any further on this observation.

### 3.2.2 Directed Graphs

In the previous section, we have addressed the representation of an independence relation in the formalism of undirected graphs. We have seen that not every independence relation can be represented faithfully in this formalism. In this section, we investigate the formalism of directed graphs (digraphs) and its expressive power for representing independence relations; we will consider acyclic digraphs only.

Just like undirected graphs, do directed graphs not have a probabilistic meaning by themselves and therefore have to be assigned one. For this purpose, we build on the same idea as outlined in the previous section for undirected graphs. We once more model the variables of an independence relation as vertices in a digraph and represent its independence statements by absence of arcs. For undirected graphs, this basic idea has been formalised in the separation criterion. For digraphs, we will now formulate a similar criterion, called the $d$-separation criterion. Before defining this criterion, however, we introduce the concept of blocking influence among variables. The definition provided here is an enhancemnet of the original definition by Pearl, based upon more recent insights [Van der Gaag \& Meyer, 1998].

Definition 3.2.9 Let $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ be an acyclic digraph and let $s$ be a chain in $G$ between $V_{i} \in \boldsymbol{V}_{G}$ and $V_{j} \in \boldsymbol{V}_{G}$. The chain $s$ is blocked by a (possibly empty) set of vertices $\boldsymbol{W} \subseteq \boldsymbol{V}_{G}$ if $V_{i} \in \boldsymbol{W}$ or $V_{j} \in \boldsymbol{W}$, or $s$ contains three consecutive vertices $X_{1}$, $X_{2}, X_{3}$, for which (at least) one of the following conditions hold:
a. arcs $\left(X_{2}, X_{1}\right)$ and $\left(X_{2}, X_{3}\right)$ are on the chain $s$, and $X_{2} \in \boldsymbol{W}$;
b. arcs $\left(X_{1}, X_{2}\right)$ and $\left(X_{2}, X_{3}\right)$ are on the chain $s$, and $X_{2} \in \boldsymbol{W}$;
c. arcs $\left(X_{1}, X_{2}\right)$ and $\left(X_{3}, X_{2}\right)$ are on the chain $s$, and $\boldsymbol{\sigma}_{G}^{*}\left(X_{2}\right) \cap \boldsymbol{W}=\varnothing$.

In defining the concept of a blocked chain, we have distinguished three graphical conditions. Figure 3.4 serves as a reference for these conditions; in the two chains representing the conditions $a$. and $b$., vertex $X_{2}$ is drawn with shading to indicate that it is comprised in the blocking set $\boldsymbol{W}$ for the chain at hand.

Condition $a$.


Condition $b$.


Condition c.


Figure 3.4: Chain Blocking; Shaded Vertices are Members of Blocking-set $\boldsymbol{W}$.
Building on the concept of blocking of single chains, we now define the d-separation criterion for reading from a digraph sets of vertices that allow for blocking all chains between two given sets of vertices.

Definition 3.2.10 Let $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ be an acyclic digraph. Let $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}_{G}$ be sets of vertices in $G$. The set of vertices $\boldsymbol{Z}$ is said to d-separate the sets of vertices $\boldsymbol{X}$ and $\boldsymbol{Y}$ in $G$, denoted as $\langle\boldsymbol{X}| \boldsymbol{Z}|\boldsymbol{Y}\rangle_{G}^{d}$, if for each vertex $V_{i} \in \boldsymbol{X}$ and each vertex $V_{j} \in \boldsymbol{Y}$ every chain from $V_{i}$ to $V_{j}$ in $G$ is blocked by $\boldsymbol{Z}$.

We would like to note that in assigning a meaning to the topological properties of a directed graph, we want to distinguish between conditional independences and dependences, that is, between two alternatives only. The formalism of directed graphs, however, allows for distinguishing between three alternatives since there are three different ways in which two arcs between three vertices can be directed (up to renaming of vertices). This observation accounts for two conditions of the concept of blocking having been assigned the same meaning; these are the first two conditions depicted in Figure 3.4. Note that the third condition models an induced dependence.

The following definition now relates the d-separation criterion to the concept of independence.

Definition 3.2.11 Let $\boldsymbol{V}$ be a set of statistical variables and let $I$ be an independence relation on $\boldsymbol{V}$. Furthermore, let $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ be an acyclic digraph with $\boldsymbol{V}_{G}=\boldsymbol{V}$. Then,


Figure 3.5: An Example Directed D-map.

- the graph $G$ is called a directed dependence map, or D-map for short, for I, if for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}$, we have: if $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y})$ then $\langle\boldsymbol{X}| \boldsymbol{Z}|\boldsymbol{Y}\rangle_{G}^{d}$;
- the graph $G$ is called a directed independence map, or I-map for short, for $I$, if for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}$, we have: if $\langle\boldsymbol{X}| \boldsymbol{Z}|\boldsymbol{Y}\rangle_{G}^{d}$ then $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y})$;
- the graph $G$ is called a directed perfect map, or P-map for short, for I, if $G$ is both a directed D-map and a directed I-map for I.

The different types of directed map have the same meaning as the different types of undirected map we discerned in Section 3.2.1. In a directed D-map for an independence relation, any pair of neighbouring vertices again represents a pair of dependent variables; however, not every pair of dependent variables needs be represented as a pair of neighbouring vertices.

Example 3.2.12 Let $\boldsymbol{V}=\left\{V_{1}, V_{2}, V_{3}, V_{3}\right\}$ be a set of statistical variables. We consider the independence relation $I$ on $\boldsymbol{V}$ that is defined by the independence statements $I\left(\left\{V_{1}\right\}, \varnothing,\left\{V_{2}\right\}\right)$ and $I\left(\left\{V_{1}, V_{2}\right\},\left\{V_{3}\right\},\left\{V_{4}\right\}\right)$. The digraph $G$ shown in Figure 3.5 is an example of a directed D-map for $I$, since for each independence statement $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y})$, $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}$, from the relation $I$, there exists a matching d-separation statement $\langle\boldsymbol{X}| \boldsymbol{Z}|\boldsymbol{Y}\rangle_{G}^{d}$. For example, for the independence statement $I\left(\left\{V_{1}\right\}, \varnothing,\left\{V_{2}\right\}\right)$, we find that $\left\langle\left\{V_{1}\right\}\right| \varnothing\left|\left\{V_{2}\right\}\right\rangle_{G}^{d}$ as there does not exist any chain between the vertices $V_{1}$ and $V_{2}$ in $G$.

In a directed I-map, any pair of non-neighbouring vertices once more represents a pair of variables that are independent; however, not every pair of independent variables is represented as a pair of non-neighbouring vertices.

Example 3.2.13 Consider once more the independence relation $I$ from Example 3.2.12. The digraph $G$ shown in Figure 3.6 is an example of a directed I-map for the relation $I$, since for each d-separation statement $\langle\boldsymbol{X}| \boldsymbol{Z}|\boldsymbol{Y}\rangle_{G}^{d}, \boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}$, read from $G$, there exists a matching independence statement $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y})$ in $I$. For example, for the d-separation statement $\left\langle\left\{V_{1}\right\}\right|\left\{V_{3}\right\}\left|\left\{V_{4}\right\}\right\rangle_{G}^{d}$ read from the digraph, we find the statement $I\left(\left\{V_{1}\right\},\left\{V_{3}\right\},\left\{V_{4}\right\}\right)$ in the independence relation.

A directed P-map for an independence relation once more faithfully captures all independences and dependences from the relation.

Example 3.2.14 Consider once more the independence relation $I$ from Example 3.2.12. The digraph $G$ shown in Figure 3.7 is a directed P-map for $G$. Note that $G$ is a directed D-map as well as a directed I-map for $I$.

From the above observations, we once more conclude that, if a digraph is a directed I-map for an independence relation, we can read independence statements from it; from a directed D-map we can read dependence statements, and from a directed P-map we can read both independence and dependence statements.

Just as we have done for undirected graphs, we investigate the expressive power of the formalism of directed graphs for representing independence relations. It can easily be shown that every independence relation has a directed D-map and a directed I-map.

Lemma 3.2.15 For every independence relation, there exist a directed $D$-map and $a$ directed I-map.

Unfortunately, a similar property does not hold for directed P-maps: not every independence relation has a directed P-map.

Example 3.2.16 Consider an independence relation $I$ on a set of variables $\boldsymbol{V}$ such that $I\left(\left\{V_{1}\right\},\left\{V_{2}, V_{3}\right\},\left\{V_{4}\right\}\right)$ and $I\left(\left\{V_{2}\right\},\left\{V_{1}, V_{4}\right\},\left\{V_{3}\right\}\right)$ for some variables $V_{1}, V_{2}, V_{3}, V_{4} \in \boldsymbol{V}$. No acyclic digraph can model these independence statements simultaneously.

Although not every independence relation has a directed P-map, there are independence relations that can indeed be represented faithfully. An independence relation for which a directed P-map exists is termed directed graph-isomorphic.

Definition 3.2.17 An independence relation $I$ is said to be directed graph-isomorphic if there exists an acyclic digraph $G$ such that $G$ is a directed $P$-map for $I$.

We would like to note that, if an independence relation is directed graph-isomorphic, it may allow several different directed P-maps, that is, a directed P-map does not need to be unique. To conclude, we would like to mention that a necessary condition has been identified for an independence relation to be directed graph-isomorphic [Pearl, 1988]. Here, we will not elaborate any further on this observation.

### 3.2.3 Choosing a Graphical Representation

In the previous sections we have investigated two different graphical formalisms for representing independence relations: the formalism of undirected graphs and the formalism of directed graphs. Both formalisms allow for a space-efficient representation of an independence relation: all graphical representations considered are polynomial in terms of the number of variables of the relation at hand. In addition, the two formalisms provide for an explicit representation of independence that allows for efficiently


Figure 3.6: An Example Directed I-map.


Figure 3.7: An Example Directed P-map.
verifying independence statements without requiring numerical computations. These properties render graphical formalisms highly suitable for representing independence relations in knowledge-based systems. Unfortunately, while the graphical formalisms allow for capturing some independence relations to accuracy, they do not allow for a faithful representation for every such relation. The correspondence between independence relations on the one hand and graphical representations on the other hand is shown schematically in Figure 3.8. This figure depicts that there are independence relations that are undirected graph-isomorphic yet not directed graph-isomorphic, and vice versa; also, there are independence relations that are both undirected and directed graph-isomorphic and independence relations that are not graph-isomorphic at all. This property negatively affects the suitability of graphical formalisms for representing independence relations. The efficiency of representation, however, generally is considered to outweigh the lack of expressive power of these formalisms. We would like to note that it is possible to mix directed and undirected graphs; such mixed graphs are called chain graphs [Studený, 1998] and are capable of representing a broader class of independence relations. Chain graphs are, however, hardly used for practical applications as a result of their more complex semantics.

In the Bayesian network framework, therefore, a graphical formalism is used for representing independences. As experience learns that the independence relations that are typically encountered in practical problem domains often are best represented by a directed graph rather than by an undirected graph, the formalism of directed graphs is employed in the framework. There does exist an undirected counterpart of the Bayesian network in the family of Probabilistic Graphical Models: the Markov network.

In real-life problem domains, independence relations may be encountered that are not graph-isomorphic. For such an independence relation, it is not possible to faithfully represent all independences as well as all dependences. The relation therefore has to be represented in either a directed I-map or a directed D-map. Now observe that since we are interested in exploiting independences for simplifying computations, we have to make sure that independences that can be read from the graphical representation of an independence relation actually do hold in the relation at hand; otherwise, we would assume independences where there are none and thereby introduce errors in inference. An independence relation that is not directed graph-isomorphic therefore is best represented by a directed I-map. Furthermore, acknowledging that some independences will escape representation, we settle for a representation in which as many of the independences of the relation as possible are modelled, that is, we insist that the number of unrepresented independences is kept at a minimum. A directed I-map that does not contain any superfluous arcs is called minimal.


Figure 3.8: Independence Relations and Graphical Representations.


Figure 3.9: An Example Minimal Directed I-map.

Definition 3.2.18 Let $\boldsymbol{V}$ be a set of statistical variables and let $I$ be an independence relation on $\boldsymbol{V}$. Furthermore, let $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ be an acyclic digraph with $\boldsymbol{V}_{G}=\boldsymbol{V}$. Then, the digraph $G$ is called a minimal directed I-map for I if $G$ is a directed I-map for $I$ and no proper subgraph of $G$ is a directed I-map for $I$.

We would like to note that an independence relation may allow several minimal I-maps, that is, a minimal directed I-map need not be unique.

Example 3.2.19 Let $\boldsymbol{V}=\left\{V_{1}, V_{2}, V_{3}, V_{4}\right\}$ be a set of statistical variables. We consider once more the independence relation $I$ on $\boldsymbol{V}$ that is defined by the statements $I\left(\left\{V_{1}\right\},\left\{V_{2}, V_{3}\right\},\left\{V_{4}\right\}\right)$ and $I\left(\left\{V_{2}\right\},\left\{V_{1}, V_{4}\right\},\left\{V_{3}\right\}\right)$. In Example 3.2.16 we have argued that the relation $I$ is not directed graph-isomorphic. The digraph $G$ shown in Figure 3.9 now is an example of a minimal directed I-map for $I$. Note that, whereas the independence statement $I\left(\left\{V_{1}\right\},\left\{V_{2}, V_{3}\right\},\left\{V_{4}\right\}\right)$ is portrayed by $G$, the statement $I\left(\left\{V_{2}\right\},\left\{V_{1}, V_{4}\right\},\left\{V_{3}\right\}\right)$ has escaped representation.

## Exercises

## Exercise 3.1

Let $\boldsymbol{V}$ be a set of statistical variables. Let $\operatorname{Pr}$ be a joint probability distribution on $\boldsymbol{V}$ and let $I_{\operatorname{Pr}}$ be its independence relation. Show that $I_{\operatorname{Pr}}$ satisfies the properties

* a. $I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \rightarrow I_{\operatorname{Pr}}(\boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{X}) ;$
* b. $I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \rightarrow I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \wedge I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{W}) ; \quad$ (quite difficult!)
${ }^{*}$ c. $I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \rightarrow I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{W}, \boldsymbol{Y})$;
d. $I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}) \wedge I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{Y}, \boldsymbol{W}) \rightarrow I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) ;$
for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W} \subseteq \boldsymbol{V}$.


## * Exercise 3.2

Let $\boldsymbol{V}$ be a set of statistical variables and let I be a semi-graphoid independence relation on $\boldsymbol{V}$. Show that

$$
I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \wedge I(\boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W}) \rightarrow I(\boldsymbol{X} \cup \boldsymbol{W}, \boldsymbol{Z}, \boldsymbol{Y})
$$

for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W} \subseteq \boldsymbol{V}$.

## * Exercise 3.3

Let $\boldsymbol{V}$ be a set of statistical variables and let I be a semi-graphoid independence relation on $\boldsymbol{V}$. Show that

$$
I(\boldsymbol{X}, \boldsymbol{Y} \cup \boldsymbol{Z}, \boldsymbol{U} \cup \boldsymbol{W}) \wedge I(\boldsymbol{Y}, \boldsymbol{Z} \cup \boldsymbol{U}, \boldsymbol{X}) \rightarrow I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{U}, \boldsymbol{Y} \cup \boldsymbol{W})
$$

for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W} \subseteq \boldsymbol{V}$.

## Exercise 3.4

Let $\boldsymbol{V}=\left\{V_{1}, V_{2}, V_{3}, V_{4}\right\}$ be a set of statistical variables. Furthermore, let $I$ be the following (semi-graphoid) independence relation on $\boldsymbol{V}$ :

| $I\left(\left\{V_{1}\right\}, \varnothing,\left\{V_{4}\right\}\right)$ | $I\left(\left\{V_{1}\right\}, \varnothing,\left\{V_{2}\right\}\right)$ | $I\left(\left\{V_{1}, V_{3}\right\},\left\{V_{2}\right\},\left\{V_{4}\right\}\right)$ |
| :--- | :--- | :--- |
| $I\left(\left\{V_{2}\right\}, \varnothing,\left\{V_{4}\right\}\right)$ | $I\left(\left\{V_{2}\right\}, \varnothing,\left\{V_{1}\right\}\right)$ | $I\left(\left\{V_{4}\right\},\left\{V_{2}\right\},\left\{V_{1}\right\}\right)$ |
| $I\left(\left\{V_{3}\right\}, \varnothing,\left\{V_{4}\right\}\right)$ | $I\left(\left\{V_{1}, V_{4}\right\}, \varnothing,\left\{V_{2}\right\}\right)$ | $I\left(\left\{V_{4}\right\},\left\{V_{2}\right\},\left\{V_{3}\right\}\right)$ |
| $I\left(\left\{V_{4}\right\}, \varnothing,\left\{V_{1}\right\}\right)$ | $I\left(\left\{V_{2}\right\}, \varnothing,\left\{V_{1}, V_{4}\right\}\right)$ | $I\left(\left\{V_{4}\right\},\left\{V_{2}\right\},\left\{V_{1}, V_{3}\right\}\right)$ |
| $I\left(\left\{V_{4}\right\}, \varnothing,\left\{V_{2}\right\}\right)$ | $I\left(\left\{V_{2}, V_{4}\right\}, \varnothing,\left\{V_{1}\right\}\right)$ | $I\left(\left\{V_{1}\right\},\left\{V_{3}\right\},\left\{V_{4}\right\}\right)$ |
| $I\left(\left\{V_{4}\right\}, \varnothing,\left\{V_{3}\right\}\right)$ | $I\left(\left\{V_{1}\right\}, \varnothing,\left\{V_{2}, V_{4}\right\}\right)$ | $I\left(\left\{V_{2}\right\},\left\{V_{3}\right\},\left\{V_{4}\right\}\right)$ |
| $I\left(\left\{V_{1}, V_{2}\right\}, \varnothing,\left\{V_{4}\right\}\right)$ | $I\left(\left\{V_{2}\right\},\left\{V_{1}\right\},\left\{V_{4}\right\}\right)$ | $I\left(\left\{V_{1}, V_{2}\right\},\left\{V_{3}\right\},\left\{V_{4}\right\}\right)$ |
| $I\left(\left\{V_{1}, V_{3}\right\}, \varnothing,\left\{V_{4}\right\}\right)$ | $I\left(\left\{V_{3}\right\},\left\{V_{1}\right\},\left\{V_{4}\right\}\right)$ | $I\left(\left\{V_{1}\right\},\left\{V_{4}\right\},\left\{V_{2}\right\}\right)$ |
| $I\left(\left\{V_{2}, V_{3}\right\}, \varnothing,\left\{V_{4}\right\}\right)$ | $I\left(\left\{V_{2}, V_{3}\right\},\left\{V_{1}\right\},\left\{V_{4}\right\}\right)$ | $I\left(\left\{V_{2}\right\},\left\{V_{4}\right\},\left\{V_{1}\right\}\right)$ |
| $I\left(\left\{V_{4}\right\}, \varnothing,\left\{V_{1}, V_{2}\right\}\right)$ | $I\left(\left\{V_{4}\right\},\left\{V_{1}\right\},\left\{V_{2}\right\}\right)$ | $I\left(\left\{V_{3}\right\},\left\{V_{1}, V_{2}\right\},\left\{V_{4}\right\}\right)$ |
| $I\left(\left\{V_{4}\right\}, \varnothing,\left\{V_{1}, V_{3}\right\}\right)$ | $I\left(\left\{V_{4}\right\},\left\{V_{1}\right\},\left\{V_{3}\right\}\right)$ | $I\left(\left\{V_{4}\right\},\left\{V_{1}, V_{2}\right\},\left\{V_{3}\right\}\right)$ |
| $I\left(\left\{V_{4}\right\}, \varnothing,\left\{V_{2}, V_{3}\right\}\right)$ | $I\left(\left\{V_{4}\right\},\left\{V_{1}\right\},\left\{V_{2}, V_{3}\right\}\right)$ | $I\left(\left\{V_{2}\right\},\left\{V_{1}, V_{3}\right\},\left\{V_{4}\right\}\right)$ |
| $I\left(\left\{V_{1}, V_{2}, V_{3}\right\}, \varnothing,\left\{V_{4}\right\}\right)$ | $I\left(\left\{V_{1}\right\},\left\{V_{2}\right\},\left\{V_{4}\right\}\right)$ | $I\left(\left\{V_{4}\right\},\left\{V_{1}, V_{3}\right\},\left\{V_{2}\right\}\right)$ |
| $I\left(\left\{V_{4}\right\}, \varnothing,\left\{V_{1}, V_{2}, V_{3}\right\}\right)$ | $I\left(\left\{V_{3}\right\},\left\{V_{2}\right\},\left\{V_{4}\right\}\right)$ | $I\left(\left\{V_{1}\right\},\left\{V_{2}, V_{3}\right\},\left\{V_{4}\right\}\right)$ |
|  |  | $I\left(\left\{V_{4}\right\},\left\{V_{2}, V_{3}\right\},\left\{V_{1}\right\}\right)$ |

Show that each statement $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y}), \boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}$, of the independence relation $I$ can be derived from the statements $I\left(\left\{V_{1}, V_{2}, V_{3}\right\}, \varnothing,\left\{V_{4}\right\}\right)$ and $I\left(\left\{V_{1}\right\}, \varnothing,\left\{V_{2}\right\}\right)$ by the independence axioms from Definition 3.1.3.

## * Exercise 3.5

Let $\boldsymbol{V}=\left\{V_{1}, V_{2}, V_{3}, V_{4}\right\}$ be a set of statistical variables. Let I be the independence relation on $\boldsymbol{V}$ that is defined by the statements $I\left(\left\{V_{1}\right\},\left\{V_{2}, V_{3}\right\},\left\{V_{4}\right\}\right)$ and $I\left(\left\{V_{2}\right\},\left\{V_{1}, V_{4}\right\},\left\{V_{3}\right\}\right)$.
a. Give all undirected D-maps for the independence relation I;
b. Give all undirected I-maps for the independence relation I.

## Exercise 3.6

Recall that $\nu_{G}$ is the neighbour set (see Definition 2.1.2). Show that for any undirected graph $G=\left(\boldsymbol{V}_{G}, \boldsymbol{E}_{G}\right)$ the following property holds: for any vertex $V_{i} \in \boldsymbol{V}_{G}$ and any vertex $V_{j} \in V_{G} \backslash\left(\left\{V_{i}\right\} \cup \nu_{G}\left(V_{i}\right)\right)$, we have that $\left\langle\left\{V_{i}\right\}\right| \nu_{G}\left(V_{i}\right)\left|\left\{V_{j}\right\}\right\rangle_{G}$.

## * Exercise 3.7

Let $G$ be the following acyclic digraph:


Examine for each of the following statements whether or not it holds in $G$ :
a. $\left\langle\left\{V_{1}\right\}\right|\left\{V_{2}, V_{3}\right\}\left|\left\{V_{6}\right\}\right\rangle_{G}^{d} ;$
b. $\left\langle\left\{V_{2}\right\}\right| \varnothing\left|\left\{V_{3}\right\}\right\rangle_{G}^{d}$;
c. $\left\langle\left\{V_{2}\right\}\right|\left\{V_{1}\right\}\left|\left\{V_{3}\right\}\right\rangle_{G}^{d}$;
d. $\left\langle\left\{V_{4}\right\}\right|\left\{V_{1}\right\}\left|\left\{V_{3}\right\}\right\rangle_{G}^{d}$;
e. $\left\langle\left\{V_{2}\right\}\right|\left\{V_{3}, V_{4}\right\}\left|\left\{V_{6}\right\}\right\rangle_{G}^{d}$;
f. $\left\langle\left\{V_{3}\right\}\right| \varnothing\left|\left\{V_{1}\right\}\right\rangle_{G}^{d}$.

## * Exercise 3.8

Let $\boldsymbol{V}=\left\{V_{1}, V_{2}, V_{3}, V_{4}\right\}$ be a set of statistical variables. Let I be the independence relation on $\boldsymbol{V}$ that is defined by the statements $I\left(\left\{V_{1}\right\}, \varnothing,\left\{V_{2}\right\}\right)$ and $I\left(\left\{V_{1}, V_{2}\right\},\left\{V_{3}\right\},\left\{V_{4}\right\}\right)$.
a. Give some directed D-maps for the independence relation I;
b. Give some directed I-maps for the independence relation I.

## * Exercise 3.9

Show that for every independence relation there exists a directed $D$-map and a directed I-map.

* Exercise 3.10

Given an example of an independence relation that has more than one directed $P$-map.

## Exercise 3.11

Let $\boldsymbol{V}=\left\{V_{1}, V_{2}, V_{3}, V_{4}\right\}$ be a set of statistical variables. Let $I$ be the independence relation on $\boldsymbol{V}$ that is defined by the statements $I\left(\left\{V_{1}\right\},\left\{V_{2}, V_{3}\right\},\left\{V_{4}\right\}\right)$ and $I\left(\left\{V_{2}\right\},\left\{V_{1}, V_{4}\right\},\left\{V_{3}\right\}\right)$. Give some minimal directed I-maps for the relation $I$.

## Exercise 3.12

Show that for any acyclic directed graph $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ the following property holds: for any vertex $V_{i} \in \boldsymbol{V}_{G}$ and any vertex $V_{j} \in \boldsymbol{V}_{G} \backslash\left(\boldsymbol{\sigma}_{G}^{*}\left(V_{i}\right) \cup \boldsymbol{\rho}_{G}\left(V_{i}\right)\right)$, we have that $\left\langle\left\{V_{i}\right\}\right| \boldsymbol{\rho}_{G}\left(V_{i}\right)\left|\left\{V_{j}\right\}\right\rangle_{G}^{d}$.

This property is called the local Markov property of a $D A G$.

## * Exercise 3.13

Let $\boldsymbol{V}$ be a set of statistical variables. Let $I$ be an independence relation on $\boldsymbol{V}$ and let $G$ be a directed I-map for $I$. Now, let $H$ be the underlying graph of $G$. Is $H$ an undirected I-map for I ?

## * Exercise 3.14

a. Give an example of an independence relation that has both an undirected $P$-map and a directed $P$-map.
b. Give an example of an independence relation that has an undirected P-map but no directed $P$-map.
c. Give an example of an independence relation that has a directed P-map but no undirected $P$-map.
d. Give an example of an independence relation that has no undirected $P$-map nor a directed P-map.

## Chapter 4

## The Bayesian Network Framework

The Bayesian network framework is characterised by a powerful and intuitively appealing formalism for representing a joint probability distribution on a set of statistical variables, for use in a knowledge-based system. Associated with this formalism are algorithms for efficiently computing probabilities of interest and for processing evidence; these algorithms constitute the basic building blocks for reasoning with knowledge represented in the formalism. In this chapter we introduce the Bayesian network formalism and detail some of the most well-known associated algorithms.

### 4.1 The Bayesian Network Formalism

The Bayesian network formalism provides a concise representation of a joint probability distribution on a set of statistical variables. The conciseness of representation is arrived at by explicit separation of knowledge of the independences holding in a distribution and the numerical quantities involved. To this end, a Bayesian network comprises two parts: a qualitative part and a quantitative part. The qualitative part of a Bayesian network is a graphical representation of the independences holding among the variables in the probability distribution that is being represented; more in specific, the qualitative part of a Bayesian network is a (minimal) directed I-map for the independence relation of the distribution. Associated with the qualitative part of a Bayesian network is a set of functions representing numerical quantities from the distribution; with each vertex in the digraph is associated a probability assessment function which basically is a set of (conditional) probabilities describing the influence of the values of the vertex' predecessors on the probabilities of the values of this vertex itself. These probability assessment functions with each other constitute the quantitative part of the Bayesian network; the numbers specified by the functions are often refered to as (network) parameters.

We define the concept of a Bayesian network more formally.
Definition 4.1.1 $A$ Bayesian network is a tuple $\mathcal{B}=(G, \Gamma)$ where

- $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is an acyclic digraph with vertices $\boldsymbol{V}_{G}=\left\{V_{1}, \ldots, V_{n}\right\}, n \geq 1$, and $\operatorname{arcs} \boldsymbol{A}_{G}$;
- $\Gamma=\left\{\gamma_{V_{i}} \mid V_{i} \in \boldsymbol{V}_{G}\right\}$ is a set of real-valued non-negative functions

$$
\gamma_{V_{i}}:\left\{c_{V_{i}}\right\} \times\left\{c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right\} \rightarrow[0,1]
$$



Figure 4.1: The Digraph of an Example Bayesian Network.
called (conditional) probability assessment functions, such that for each configuration $c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}$ of the set $\boldsymbol{\rho}_{G}\left(V_{i}\right)$ of (immediate) predecessors of vertex $V_{i}$ in $G$, we have that $\gamma_{V_{i}}\left(\neg v_{i} \mid c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)=1-\gamma_{V_{i}}\left(v_{i} \mid c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right), i=1, \ldots, n$.

Note from the previous definition that there exists a one-to-one correspondence between vertices $V_{i} \in \boldsymbol{V}_{G}$ and statistical variables $V_{i} \in \boldsymbol{V}$.

Example 4.1.2 We consider the digraph $G$ shown in Figure 4.1. With the digraph $G$, we associate a set $\Gamma=\left\{\gamma_{V_{i}} \mid i=1, \ldots, 4\right\}$ of probability assessment functions $\gamma_{V_{i}}$. For example, for the vertices $V_{1}$ and $V_{2}$, the probability assessment functions $\gamma_{V_{1}}$ and $\gamma_{V_{2}}$ may be defined as

$$
\begin{array}{lllll}
\gamma_{V_{1}}\left(v_{1}\right) & =0.25 \\
\gamma_{V_{1}}\left(\neg v_{1}\right) & =0.75 & & \text { and } & \gamma_{V_{2}}\left(v_{2}\right)
\end{array}=0.5
$$

For vertex $V_{4}$, the probability assessment function $\gamma_{V_{4}}$ is defined as

$$
\begin{array}{llll}
\gamma_{V_{4}}\left(v_{4} \mid v_{3}\right) & =0.8 & \gamma_{V_{4}}\left(\neg v_{4} \mid v_{3}\right) & =0.2 \\
\gamma_{V_{4}}\left(v_{4} \mid \neg v_{3}\right) & =0 & \gamma_{V_{4}}\left(\neg v_{4} \mid \neg v_{3}\right) & =1.0
\end{array}
$$

For vertex $V_{3}$, the probability assessment function $\gamma_{V_{3}}$ is defined by the values

$$
\begin{array}{ll}
\gamma_{V_{3}}\left(v_{3} \mid v_{1} \wedge v_{2}\right) & =0.75 \\
\gamma_{V_{3}}\left(v_{3} \mid \neg v_{1} \wedge v_{2}\right) & =0.4 \\
\gamma_{V_{3}}\left(v_{3} \mid v_{1} \wedge \neg v_{2}\right) & =0.25 \\
\gamma_{V_{3}}\left(v_{3} \mid \neg v_{1} \wedge \neg v_{2}\right) & =0.2
\end{array}
$$

and their complementary values

$$
\begin{array}{ll}
\gamma_{V_{3}}\left(\neg v_{3} \mid v_{1} \wedge v_{2}\right) & =0.25 \\
\gamma_{V_{3}}\left(\neg v_{3} \mid \neg v_{1} \wedge v_{2}\right) & =0.6 \\
\gamma_{V_{3}}\left(\neg v_{3} \mid v_{1} \wedge \neg v_{2}\right) & =0.75 \\
\gamma_{V_{3}}\left(\neg v_{3} \mid \neg v_{1} \wedge \neg v_{2}\right) & =0.8
\end{array}
$$

The pair $\mathcal{B}=(G, \Gamma)$ is a Bayesian network.
The probability assessment functions of a Bayesian network with each other provide all information necessary for uniquely defining a joint probability distribution on the variables discerned that respects the independence relation portrayed by the qualitative part of the network; henceforth, we will call this distribution the joint probability distribution defined by the network.

Proposition 4.1.3 Let $\mathcal{B}=(G, \Gamma)$ be a Bayesian network. Then,

$$
\operatorname{Pr}\left(\boldsymbol{V}_{G}\right)=\prod_{V_{i} \in \boldsymbol{V}_{G}} \gamma_{V_{i}}\left(V_{i} \mid \boldsymbol{\rho}_{G}\left(V_{i}\right)\right)
$$

defines a joint probability distribution $\operatorname{Pr}$ on $\boldsymbol{V}_{G}$ such that $G$ is a directed I-map for the independence relation $I_{\mathrm{Pr}}$ of Pr .

Proof. Since the digraph $G$ of the Bayesian network $B$ is acyclic, it allows a total ordering of its vertices such that any successor of a vertex in the digraph follows it in the ordering; such an ordering is termed a topological order of the digraph's vertices. Note that any topological order of the vertices of $G$ constitutes an ordering of the corresponding statistical variables. We take $\iota_{G}$ to be such a topological order; for ease of exposition, we assume that $\iota_{G}\left(V_{i}\right)=i$ for all $V_{i} \in \boldsymbol{V}_{G}$. We now consider an arbitrary joint probability distribution $P$ on $\boldsymbol{V}_{G}$ such that $G$ is a directed I-map for the independence relation of $P$. To the expression $P\left(\boldsymbol{V}_{G}\right)$ describing the distribution $P$, we apply the chain rule from probability theory, such that every variable $V_{i}$ from $\boldsymbol{V}_{G}$ is conditioned on the variables $V_{1}, \ldots, V_{i-1}$, preceding it in the ordering $\iota_{G}$, that is,

$$
P\left(\boldsymbol{V}_{G}\right)=\prod_{V_{i} \in \boldsymbol{V}_{G}} P\left(V_{i} \mid V_{1} \wedge \cdots \wedge V_{i-1}\right)
$$

From the digraph $G$, we read, by means of the d-separation criterion, that $\left\langle\left\{V_{i}\right\}\right| \boldsymbol{\rho}_{G}\left(V_{i}\right) \mid$ $\left.\left\{V_{1}, \ldots, V_{i-1}\right\} \backslash \boldsymbol{\rho}_{G}\left(V_{i}\right)\right\rangle_{G}^{d}$ for every vertex $V_{i} \in \boldsymbol{V}_{G}$. Since $G$ is a directed I-map for the distribution $P$, we have for every $V_{i}$ that $I_{P}\left(\left\{V_{i}\right\}, \boldsymbol{\rho}_{G}\left(V_{i}\right),\left\{V_{1}, \ldots, V_{i-1}\right\} \backslash \boldsymbol{\rho}_{G}\left(V_{i}\right)\right)$. From this observation, we have that

$$
P\left(V_{i} \mid V_{1} \wedge \cdots \wedge V_{i-1}\right)=P\left(V_{i} \mid \boldsymbol{\rho}_{G}\left(V_{i}\right)\right)
$$

for every vertex $V_{i} \in \boldsymbol{V}_{G}$. By exploiting this property, we find that

$$
P\left(\boldsymbol{V}_{G}\right)=\prod_{V_{i} \in \boldsymbol{V}_{G}} P\left(V_{i} \mid \boldsymbol{\rho}_{G}\left(V_{i}\right)\right)
$$

Since $P$ has been chosen arbitrarily, we conclude that any joint probability distribution such that $G$ is a directed I-map for its independence relation, satisfies this property. By reversing the argument, we find that there exists a (unique) joint probability distribution $\operatorname{Pr}$ on $\boldsymbol{V}_{G}$ such that $G$ is a directed I-map for the independence relation of Pr and $\operatorname{Pr}\left(V_{i} \mid \boldsymbol{\rho}_{G}\left(V_{i}\right)\right)=\gamma_{V_{i}}\left(V_{i} \mid \boldsymbol{\rho}_{G}\left(V_{i}\right)\right)$ for each variable $V_{i} \in \boldsymbol{V}_{G}$.

We illustrate the basic idea of the previous proposition by means of an example.
Example 4.1.4 We consider once more the Bayesian network $\mathcal{B}=(G, \Gamma)$ from Example 4.1.2. From the property stated in Proposition 4.1.3, we have that from the expression

$$
\operatorname{Pr}\left(V_{1} \wedge V_{2} \wedge V_{3} \wedge V_{4}\right)=\gamma_{V_{1}}\left(V_{1}\right) \cdot \gamma_{V_{2}}\left(V_{2}\right) \cdot \gamma_{V_{3}}\left(V_{3} \mid V_{1} \wedge V_{2}\right) \cdot \gamma_{V_{4}}\left(V_{4} \mid V_{3}\right)
$$

any probability of interest can be computed. For example, we have that

$$
\begin{aligned}
\operatorname{Pr}\left(v_{1} \wedge \neg v_{2} \wedge v_{3} \wedge v_{4}\right) & =\gamma_{V_{1}}\left(v_{1}\right) \cdot \gamma_{V_{2}}\left(\neg v_{2}\right) \cdot \gamma_{V_{3}}\left(v_{3} \mid v_{1} \wedge \neg v_{2}\right) \cdot \gamma_{V_{4}}\left(v_{4} \mid v_{3}\right)= \\
& =0.25 \cdot 0.5 \cdot 0.25 \cdot 0.8=0.025
\end{aligned}
$$

and

$$
\begin{aligned}
\operatorname{Pr}\left(\neg v_{1} \wedge \neg v_{2} \wedge v_{3} \wedge v_{4}\right) & =\gamma_{V_{1}}\left(\neg v_{1}\right) \cdot \gamma_{V_{2}}\left(\neg v_{2}\right) \cdot \gamma_{V_{3}}\left(v_{3} \mid \neg v_{1} \wedge \neg v_{2}\right) \cdot \gamma_{V_{4}}\left(v_{4} \mid v_{3}\right)= \\
& =0.06
\end{aligned}
$$

By marginalisation we can now compute the probability $\operatorname{Pr}\left(\neg v_{2} \wedge v_{3} \wedge v_{4}\right)$ :

$$
\begin{aligned}
\operatorname{Pr}\left(\neg v_{2} \wedge v_{3} \wedge v_{4}\right) & =\operatorname{Pr}\left(v_{1} \wedge \neg v_{2} \wedge v_{3} \wedge v_{4}\right)+\operatorname{Pr}\left(\neg v_{1} \wedge \neg v_{2} \wedge v_{3} \wedge v_{4}\right)= \\
& =0.085
\end{aligned}
$$

From Proposition 4.1.3 it is readily seen that the function values of a Bayesian network's probability assessment functions can be interpreted as conditional probabilities.

Corollary 4.1.5 Let $\mathcal{B}=(G, \Gamma)$ be a Bayesian network and let $\operatorname{Pr}$ be the joint probability distribution defined by $\mathcal{B}$. Then, for each vertex $V_{i} \in \boldsymbol{V}_{G}$, we have that

$$
\operatorname{Pr}\left(V_{i} \mid \boldsymbol{\rho}_{G}\left(V_{i}\right)\right)=\gamma_{V_{i}}\left(V_{i} \mid \boldsymbol{\rho}_{G}\left(V_{i}\right)\right)
$$

### 4.2 Probabilistic Inference

A Bayesian network generally is used for probabilistic inference, that is, for making probabilistic statements concerning the variables that are represented in the network. For this purpose, the property stated in Proposition 4.1.3 can be exploited: from the proposition we have that, since the digraph of a Bayesian network and its associated probability assessment functions with each other uniquely define a joint probability distribution, any (prior or posterior) probability of interest can be computed from the network. Explicitly generating the represented probability distribution as indicated by the proposition and then using the basic rules of marginalisation and conditioning, however, is computationally infeasible. In addition, such a straightforward approach would not exploit the independences that are represented by the qualitative part of the network at hand. More efficient algorithms for probabilistic inference have been designed that do exploit the represented independences. The best known of these are the algorithm of J. Pearl [Pearl, 1988] and the algorithm of S.L. Lauritzen and D.J. Spiegelhalter [Lauritzen \& Spiegelhalter, 1988]. In the sequel, we will focus on Pearl's algorithm.

The basic idea of Pearl's algorithm for probabilistic inference is best explained from an object-centered point of view. The digraph of a Bayesian network is looked upon as a computational architecture: the vertices of the digraph are autonomous objects and its arcs are bi-directional communication channels. Each vertex has a local processor that is capable of performing simple probabilistic computations and a local memory in which its associated probability assessment function is stored. Through the communication channels the vertices send each other parameters ${ }^{1}$ providing information about the joint probability distribution that is represented by the network - that is, information determined from the network's probability assessments and the independences portrayed in the digraph - and about the evidence that has been entered and

[^3]processed so far. Each vertex is now able to compute the probabilities of its values from its own probability assessment function and the information it receives from its neighbours.

Initially, a Bayesian network is in an equilibrium state: recomputation of the various parameters that the vertices send one another, does not result in a change in any of them. Now, when a piece of evidence is entered into the network for some vertex, this equilibrium is perturbed, since the probability distribution over the network's variables under consideration has now changed: the distribution should be conditioned on the (additional) evidence obtained. Once the value of a variable is known with certainty, the probability distribution for the variable degenerates and can no longer change; entering evidence for a variable into a Bayesian network thus basically amounts to adding its corresponding vertex to the (initially empty) blocking-set $W$ for the digraph, thereby changing the independences that have to be taken into account by the inference algorithm. Therefore, to process the evidence entered for some vertex, the parameters this vertex sends to its neighbours are updated. After receiving updated parameters, these neighbours in turn compute new parameters to send to their neighbours, and so on. The impact of the evidence thus spreads throughout the network by messagepassing between neighbouring vertices.

In this section, we detail the computations involved in Pearl's algorithm. In doing so, we distinguish between various types of graph. In Section 4.2.1, we focus on directed trees. In Section 4.2.2 we address singly connected digraphs more in general. In Section 4.2 .3 we discuss probabilistic inference with Bayesian networks comprising a multiply connected digraph.

### 4.2.1 Directed Trees

In discussing Pearl's algorithm, we begin by focusing on probabilistic inference with a Bayesian network comprising a directed tree for its qualitative part, that is, in the network's digraph a vertex may have several successors but at most one predecessor. Before detailing the computations involved in inference with such a network, we introduce some more terminology and notational convention that we will use in the sequel. Let $G$ be the digraph of a Bayesian network, and let $\boldsymbol{V}$ be (a subset of) the set of variables represented by its vertices. Now $V_{i} \in \boldsymbol{V}$ is called instantiated if its value is known with certainty; otherwise, it is called uninstantiated. Now, let $\boldsymbol{X} \subseteq \boldsymbol{V}$ be the set of instantiated variables or vertices from $\boldsymbol{V}$. The configuration $c_{\boldsymbol{X}}$ of $\boldsymbol{X}$ that is known with certainty is called a partial configuration of $\boldsymbol{V}$ and will be denoted as $\widetilde{c}_{\boldsymbol{V}}$. Note that the notation $\widetilde{c}_{V}$ provides for referring to the subset of instantiated variables or vertices in a Bayesian network without having to specify this subset explicitly.

At any time during probabilistic inference with a Bayesian network, the probabilities of the values of a variable of interest are dependent upon all evidence entered so far into the network.

Lemma 4.2.1 Let $\mathcal{B}=(G, \Gamma)$ be a Bayesian network where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is a directed tree, and let $\operatorname{Pr}$ be the joint probability distribution defined by $\mathcal{B}$. Let $V_{i} \in \boldsymbol{V}_{G}$ be a vertex in $G$, and let $\boldsymbol{V}_{\boldsymbol{i}}^{-}=\boldsymbol{\sigma}_{G}^{*}\left(V_{i}\right)$ and $\boldsymbol{V}_{\boldsymbol{i}}^{+}=\boldsymbol{V}_{G} \backslash \boldsymbol{V}_{\boldsymbol{i}}^{-}$. Then,

$$
\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{G}}\right)=\alpha \cdot \operatorname{Pr}\left(\widetilde{c}_{V_{i}^{-}} \mid V_{i}\right) \cdot \operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{i}^{+}}\right)
$$

where $\widetilde{c}_{\boldsymbol{V}_{G}}=\widetilde{c}_{\boldsymbol{V}_{i}} \wedge \widetilde{c}_{\boldsymbol{V}_{i}}{ }^{+}$and $\alpha$ is a normalisation constant.


Figure 4.2: A Fragment of a Directed Tree.

Proof. For the probabilities $\operatorname{Pr}\left(V_{i} \mid \widetilde{C}_{V_{G}}\right)$ of the values of vertex $V_{i}$, we have

$$
\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{G}}\right)=\frac{\operatorname{Pr}\left(\widetilde{c}_{V_{G}} \mid V_{i}\right) \cdot \operatorname{Pr}\left(V_{i}\right)}{\operatorname{Pr}\left(\widetilde{c}_{V_{G}}\right)}=\frac{\operatorname{Pr}\left(\widetilde{c}_{V_{i}} \wedge \widetilde{c}_{V_{i}^{+}} \mid V_{i}\right) \cdot \operatorname{Pr}\left(V_{i}\right)}{\operatorname{Pr}\left(\widetilde{c}_{V_{i}^{-}} \wedge \widetilde{c}_{V_{i}+}\right)}
$$

by Bayes' Theorem. Now consider Figure 4.2 showing a fragment of the directed tree $G$ of the network. Using the d-separation criterion, we observe that $\langle\boldsymbol{X}|\left\{V_{i}\right\}|\boldsymbol{Y}\rangle_{G}^{d}$ for all sets of vertices $\boldsymbol{X} \subseteq \boldsymbol{V}_{\boldsymbol{i}}^{-}$and $Y \subseteq \boldsymbol{V}_{\boldsymbol{i}}^{+}$. Since $G$ is a directed I-map for the joint probability distribution $\operatorname{Pr}$, we conclude that $I_{\operatorname{Pr}}\left(\boldsymbol{X},\left\{V_{i}\right\}, \boldsymbol{Y}\right)$ for all sets $\boldsymbol{X} \subseteq \boldsymbol{V}_{\boldsymbol{i}}^{-}$, $\boldsymbol{Y} \subseteq \boldsymbol{V}_{\boldsymbol{i}}^{+}$. Exploiting this observation, we find

$$
\begin{aligned}
\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{G}}\right) & =\frac{\operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{i}^{-}} \mid V_{i}\right) \cdot \operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{i}^{+}} \mid V_{i}\right) \cdot \operatorname{Pr}\left(V_{i}\right)}{\operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{i}^{-}} \wedge \widetilde{c}_{\boldsymbol{V}_{i}^{+}}\right)}= \\
& =\frac{\operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{i}^{-}} \mid V_{i}\right) \cdot \operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{\boldsymbol{V}_{i}^{+}}\right)}{\operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{i}^{-}} \mid \widetilde{c}_{\boldsymbol{V}_{i}}\right)}
\end{aligned}
$$

Now observe that the factor $\left(\operatorname{Pr}\left(\widetilde{c}_{V_{i}^{-}} \mid \widetilde{c}_{V_{i}^{+}}\right)\right)^{-1}$ in the expression depends on the location of vertex $V_{i}$ in the digraph and, if instantiated, on the value of its instantiation, but not on the value of $V_{i}$ under consideration. It therefore is a constant with respect to $V_{i}$. In the sequel, this constant will universally be denoted as $\alpha$. The constant $\alpha$ is generally referred to as a normalisation constant because it can be computed from $\operatorname{Pr}\left(v_{i} \mid \widetilde{c}_{V_{G}}\right)+\operatorname{Pr}\left(\neg v_{i} \mid \widetilde{c}_{V_{G}}\right)=1$. The property stated in the lemma now follows by substitution.

The previous lemma shows that the probabilities of the values of a vertex of interest can be expressed in terms of two factors describing the influence of evidence entered for this vertex' descendants and for all other vertices, separately. The following definition introduces some new terminology for these separate factors.

Definition 4.2.2 Let $\mathcal{B}=(G, \Gamma)$ be a Bayesian network where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is a directed tree, and let $\operatorname{Pr}$ be the joint probability distribution defined by $\mathcal{B}$. Let $V_{i} \in \boldsymbol{V}_{G}$
be a vertex in $G$, and let $\boldsymbol{V}_{\boldsymbol{i}}^{-}$and $\boldsymbol{V}_{\boldsymbol{i}}^{+}$be as before. The compound causal parameter $\pi_{V_{i}}$ for vertex $V_{i}$ is the function $\pi_{V_{i}}:\left\{V_{i}\right\} \rightarrow[0,1]$ defined by

$$
\pi_{V_{i}}\left(V_{i}\right)=\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{i}^{+}}\right)
$$

The compound diagnostic parameter $\lambda_{V_{i}}$ for $V_{i}$ is the function $\lambda_{V_{i}}:\left\{V_{i}\right\} \rightarrow[0,1]$ defined by

$$
\lambda_{V_{i}}\left(V_{i}\right)=\operatorname{Pr}\left(\widetilde{c}_{V_{i}^{-}} \mid V_{i}\right)
$$

Note that the compound diagnostic parameter for a vertex describes the combined influence on this vertex' probabilities of all evidence that has been entered for its descendants; the compound causal parameter for the vertex describes the combined influence of evidence entered for all other vertices in the digraph.

We take a closer look at the previous definition for some special cases. We begin by addressing uninstantiated vertices having either no incoming or no outgoing arcs. A directed tree has one vertex $W$ without any incoming arcs; this vertex is the root of the tree. We observe that for $W$ the set $\boldsymbol{W}^{+}$is empty. So, $\widetilde{c}_{\boldsymbol{W}^{+}}=$True. The compound causal parameter $\pi_{W}$ for $W$ therefore equals $\pi_{W}(W)=\operatorname{Pr}(W)$. The directed tree may further include several vertices having no outgoing arcs; these vertices are the leaves of the tree. For a leaf $V$, we observe that the set $\boldsymbol{V}^{-}$consists of $V$ only. From $V$ being uninstantiated, we have that $\widetilde{c}_{V^{-}}=$True. The compound diagnostic parameter $\lambda_{V}$ for $V$ therefore equals $\lambda_{V}(V)=1$. To conclude, we consider instantiated vertices. For a vertex $V_{i}$ for which the evidence $V_{i}=$ true has been entered, we find $\pi_{V_{i}}\left(v_{i}\right)=\operatorname{Pr}\left(v_{i} \mid \widetilde{c}_{V_{i}^{+}}\right)$and $\pi_{V_{i}}\left(\neg v_{i}\right)=\operatorname{Pr}\left(\neg v_{i} \mid \widetilde{c}_{V_{i}^{+}}\right)$, and $\lambda_{V_{i}}\left(v_{i}\right)=\operatorname{Pr}\left(\widetilde{c}_{V_{i}^{-}} \mid v_{i}\right)$ and $\lambda_{V_{i}}\left(\neg v_{i}\right)=0$; an analogous observation holds for the case where the evidence $V_{i}=$ false has been entered.

Using the definition of the compound causal and diagnostic parameters for a vertex, the property stated in Lemma 4.2 .1 can now be reformulated: for each vertex $V_{i}$, we have that

$$
\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{G}}\right)=\alpha \cdot \pi_{V_{i}}\left(V_{i}\right) \cdot \lambda_{V_{i}}\left(V_{i}\right)
$$

where $\alpha$ is a normalisation constant. In this form, the lemma is known as the data fusion lemma [Pearl, 1988]. Note that the data fusion lemma implies that the compound and diagnostic parameters for a vertex provide it with enough information for computing the probabilities of its values, that is, no further knowledge of the joint probability distribution is needed.

The compound diagnostic parameter for a vertex specifies probabilistic information from all its descendants combined; an analogous observation applies to the compound causal parameter for the vertex. To be able to exploit the digraph of a Bayesian network as a computational architecture as outlined before, the compound parameters for a vertex have to be computed from separate parameters originating from its various neighbours. The following definition introduces such separate parameters; the Lemmas 4.2.4 and 4.2.5 will show the computation of the compound parameters from these separate ones.

Definition 4.2.3 Let $\mathcal{B}=(G, \Gamma)$ be a Bayesian network where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is a directed tree, and let $\operatorname{Pr}$ be the joint probability distribution defined by $\mathcal{B}$. For each
vertex $V \in \boldsymbol{V}_{G}$, let $\boldsymbol{V}^{-}$and $\boldsymbol{V}^{+}$be as before. Now, let $V_{i}$ be a vertex with a successor $V_{k}$ in $G$. The causal parameter $\pi_{V_{k}}^{V_{i}}$ from $V_{i}$ to $V_{k}$ is the function $\pi_{V_{k}}^{V_{i}}:\left\{V_{i}\right\} \rightarrow[0,1]$ defined by

$$
\pi_{V_{k}}^{V_{i}}\left(V_{i}\right)=\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{k}^{+}}\right)
$$

Now, let $V_{i}$ be a vertex having the predecessor $V_{j}$ in $G$. The diagnostic parameter $\lambda_{V_{i}}^{V_{j}}$ from $V_{i}$ to $V_{j}$ is the function $\lambda_{V_{i}}^{V_{j}}:\left\{\boldsymbol{V}_{j}\right\} \rightarrow[0,1]$ defined by

$$
\lambda_{V_{i}}^{V_{j}}\left(V_{j}\right)=\operatorname{Pr}\left(\widetilde{c}_{V_{i}^{-}} \mid V_{j}\right)
$$

Note that a causal parameter is a parameter that a vertex sends to a successor to provide this successor with information concerning its non-descendants. A diagnostic parameter is a parameter that a vertex sends to its predecessor to provide this predecessor with information concerning the vertices located in the subtree rooted at the vertex at hand. The separate causal and diagnostic parameters are the messages the vertices send each other through the communication channels of the computational architecture and, hence, may be looked upon as associated with the arcs of the directed tree of the Bayesian network at hand.

We take a closer look at the previous definition for some special cases. We observe that for the root of a directed tree no diagnostic parameter is defined because it does not have a predecessor. For the leaves of the tree no causal parameters are defined as these vertices do not have any successors. In addition, we note that for a vertex $V_{i}$ for which the evidence $V_{i}=$ true is observed and entered into the network, we find that $\pi_{V_{k}}^{V_{i}}\left(v_{i}\right)=1$ and $\pi_{V_{k}}^{V_{i}}\left(\neg v_{i}\right)=0$ for any successor $V_{k}$ of $V_{i}$; an analogous observation holds for the case where $V_{i}=$ false is observed.

The following lemma now shows that a vertex can compute its compound causal parameter from the causal parameter it receives from its predecessor and its own probability assessment function.

Lemma 4.2.4 Let $\mathcal{B}=(G, \Gamma)$ be a Bayesian network where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is a directed tree. Let $V_{i} \in V_{G}$ be a vertex with the predecessor $V_{j}$ in $G$. Let $\pi_{V_{i}}$ be the compound causal parameter for vertex $V_{i}$, and let $\pi_{V_{i}}^{V_{j}}$ be the causal parameter from $V_{j}$ to $V_{i}$. Then,

$$
\pi_{V_{i}}\left(V_{i}\right)=\sum_{c_{V_{j}}} \gamma_{V_{i}}\left(V_{i} \mid c_{V_{j}}\right) \cdot \pi_{V_{i}}^{V_{j}}\left(c_{V_{j}}\right)
$$

Proof. Let Pr be the joint probability distribution defined by the Bayesian network $\mathcal{B}$. For vertex $V_{i}$, let $\boldsymbol{V}_{\boldsymbol{i}}^{+}$be as before. Then, by definition we have that

$$
\pi_{V_{i}}\left(V_{i}\right)=\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{\boldsymbol{V}_{i}^{+}}\right)
$$

By conditioning on the values of $V_{i}$ 's predecessor $V_{j}$, we find

$$
\pi_{V_{i}}\left(V_{i}\right)=\operatorname{Pr}\left(V_{i} \mid v_{j} \wedge \widetilde{c}_{\boldsymbol{V}_{\boldsymbol{i}}^{+}}\right) \cdot \operatorname{Pr}\left(v_{j} \mid \widetilde{c}_{\boldsymbol{V}_{\boldsymbol{i}}^{+}}\right)+\operatorname{Pr}\left(V_{i} \mid \neg v_{j} \wedge \widetilde{c}_{\boldsymbol{V}_{\boldsymbol{i}}^{+}}\right) \cdot \operatorname{Pr}\left(\neg v_{j} \mid \widetilde{c}_{\boldsymbol{V}_{\boldsymbol{i}}^{+}}\right)
$$

Now consider once more Figure 4.2 showing a fragment of the directed tree $G$ of the network. Using the d-separation criterion, we observe that $\left\langle\left\{V_{i}\right\}\right|\left\{V_{j}\right\}|\boldsymbol{X}\rangle_{G}^{d}$ for all sets
of vertices $\boldsymbol{X} \subseteq \boldsymbol{V}_{\boldsymbol{i}}^{+}$. Since $G$ is a directed I-map for the distribution Pr, we have that $I_{\operatorname{Pr}}\left(\left\{V_{i}\right\},\left\{V_{j}\right\}, \boldsymbol{X}\right)$ for all sets $\boldsymbol{X} \subseteq \boldsymbol{V}_{\boldsymbol{i}}^{+}$. Exploiting this observation, we find

$$
\pi_{V_{i}}\left(V_{i}\right)=\operatorname{Pr}\left(V_{i} \mid v_{j}\right) \cdot \operatorname{Pr}\left(v_{j} \mid \widetilde{c}_{V_{i}^{+}}\right)+\operatorname{Pr}\left(V_{i} \mid \neg v_{j}\right) \cdot \operatorname{Pr}\left(\neg v_{j} \mid \widetilde{c}_{V_{i}^{+}}\right)
$$

The probabilities $\operatorname{Pr}\left(V_{i} \mid v_{j}\right)$ and $\operatorname{Pr}\left(V_{i} \mid \neg v_{j}\right)$ have been specified as function values of the probability assessment function $\gamma_{V_{i}}$ and therefore are directly available to vertex $V_{i}$. In addition, vertex $V_{i}$ receives the probabilities $\operatorname{Pr}\left(v_{j} \mid \widetilde{c}_{V_{i}^{+}}\right)$and $\operatorname{Pr}\left(\neg v_{j} \mid \widetilde{c}_{V_{i}^{+}}\right)$ from its predecessor $V_{j}$ as function values of the causal parameter $\pi_{V_{i}}^{V_{j}}$. The property stated in the lemma now follows by substitution.

A vertex can further compute its compound diagnostic parameter from the separate diagnostic parameters it receives from its successors in the digraph. This property is stated more formally in the following lemma.

Lemma 4.2.5 Let $\mathcal{B}=(G, \Gamma)$ be a Bayesian network where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is a directed tree. Let $V_{i} \in V_{G}$ be an uninstantiated vertex with the successors $\boldsymbol{\sigma}_{G}\left(V_{i}\right)=$ $\left\{V_{i_{1}}, \ldots, V_{i_{m}}\right\}, m \geq 1$, in $G$. Furthermore, let $\lambda_{V_{i}}$ be the compound diagnostic parameter for vertex $V_{i}$ and, for each $V_{i_{j}} \in \sigma_{G}\left(V_{i}\right)$, let $\lambda_{V_{i_{j}}}^{V_{i}}$ be the diagnostic parameter from $V_{i_{j}}$ to $V_{i}$. Then,

$$
\lambda_{V_{i}}\left(V_{i}\right)=\prod_{j=1, \ldots, m} \lambda_{V_{i_{j}}}^{V_{i}}\left(V_{i}\right)
$$

Proof. Let Pr be the joint probability distribution defined by the Bayesian network $B$. For each vertex $V \in \boldsymbol{V}_{G}$, let $\boldsymbol{V}^{-}$be as before. Since $V_{i}$ is an uninstantiated vertex, we have that $\widetilde{c}_{V_{i}^{-}}=\widetilde{c}_{V_{i_{1}}} \wedge \cdots \wedge \widetilde{c}_{V_{i_{m}}^{-}}$. So,

$$
\lambda_{V_{i}}\left(V_{i}\right)=\operatorname{Pr}\left(\widetilde{c}_{V_{i}} \mid V_{i}\right)=\operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{i_{1}}^{-}} \wedge \cdots \wedge \widetilde{c}_{\boldsymbol{V}_{i_{m}}^{-}} \mid V_{i}\right)
$$

Now consider Figure 4.3 showing a fragment of the directed tree $G$ of the network. We


Figure 4.3: Exploiting d-Separation for Computing $\lambda_{V_{i}}\left(V_{i}\right)$.
observe that $\langle\boldsymbol{X}|\left\{V_{i}\right\}|\boldsymbol{Y}\rangle_{G}^{d}$ for all sets of vertices $\boldsymbol{X} \subseteq \boldsymbol{V}_{i_{j}}^{-}$and $\boldsymbol{Y} \subseteq \bigcup_{k=1, \ldots, m, k \neq j} \boldsymbol{V}_{i_{k}}^{-}$, $j=1, \ldots, m$. Since $G$ is a directed I-map for the distribution Pr, we have that
$I_{\operatorname{Pr}}\left(\boldsymbol{X},\left\{V_{i}\right\}, \boldsymbol{Y}\right)$ for all sets $\boldsymbol{X} \subseteq \boldsymbol{V}_{\boldsymbol{i}_{\boldsymbol{j}}}^{-}, \boldsymbol{Y} \subseteq \bigcup_{k=1, \ldots, m, k \neq j} \boldsymbol{V}_{\boldsymbol{i}_{\boldsymbol{k}}}^{-}, j=1, \ldots, m$. It follows that

$$
\lambda_{V_{i}}\left(V_{i}\right)=\operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{i_{\mathbf{1}}}^{-}} \mid V_{i}\right) \cdot \ldots \cdot \operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{i_{m}}^{-}} \mid V_{i}\right)
$$

The probabilities $\operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{\boldsymbol{i}_{\boldsymbol{j}}}} \mid V_{i}\right)$ are sent to vertex $V_{i}$ by its successor $V_{i_{j}}$ as function values of the diagnostic parameter $\lambda_{V_{i_{j}}}^{V_{i}}, j=1, \ldots, m$. The property stated in the lemma now follows by substitution.

Note that the previous lemma applies to uninstantiated vertices only. However, the property mentioned in the lemma can be taken to hold for an instantiated vertex $V_{i}$ as well, if entering evidence for $V_{i}$ into the network is modelled by adding a 'dummy' successor $D$ for vertex $V_{i}$ that sends an appropriate diagnostic parameter to $V_{i}$. For the evidence $V_{i}=t r u e$, this 'dummy' successor sends the diagnostic parameter $\lambda_{D}^{V_{i}}$ with $\lambda_{D}^{V_{i}}\left(v_{i}\right)=1$ and $\lambda_{D}^{V_{i}}\left(\neg v_{i}\right)=0$ to $V_{i}$; an analogous observation holds for the evidence $V_{i}=$ false.

So far, we have shown that a vertex can compute the probabilities of its values from its own probability assessment function and the causal and diagnostic parameters it receives from its neighbours. Now observe that this vertex in turn has to compute parameters to send to its neighbours. The following lemma shows that a vertex can compute the diagnostic parameter to send to its predecessor from its own probability assessment function and the diagnostic parameters it receives from its successors. In other words, for this purpose it combines its own information about the joint probability distribution with the information it receives concerning the evidence entered so far for its descendants.

Lemma 4.2.6 Let $\mathcal{B}=(G, \Gamma)$ be a Bayesian network where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is a directed tree. Let $V_{i} \in V_{G}$ be a vertex with the predecessor $V_{j}$ in $G$. Let $\lambda_{V_{i}}$ be the compound diagnostic parameter for vertex $V_{i}$ and let $\lambda_{V_{i}}^{V_{j}}$ be the diagnostic parameter from $V_{i}$ to $V_{j}$. Then,

$$
\lambda_{V_{i}}^{V_{j}}\left(V_{j}\right)=\sum_{c_{V_{i}}} \lambda_{V_{i}}\left(c_{V_{i}}\right) \cdot \gamma_{V_{i}}\left(c_{V_{i}} \mid V_{j}\right)
$$

Proof. Let Pr be the joint probability distribution defined by the Bayesian network $\mathcal{B}$. For vertex $V_{i}$, let $\boldsymbol{V}_{\boldsymbol{i}}^{-}$be as before. Then, by definition we have that

$$
\lambda_{V_{i}}^{V_{j}}\left(V_{j}\right)=\operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{i}^{-}} \mid V_{j}\right)
$$

By conditioning on the values of $V_{i}$, we find

$$
\lambda_{V_{i}}^{V_{j}}\left(V_{j}\right)=\operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{\boldsymbol{i}}^{-}} \mid v_{i} \wedge V_{j}\right) \cdot \operatorname{Pr}\left(v_{i} \mid V_{j}\right)+\operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{\boldsymbol{i}}^{-}} \mid \neg v_{i} \wedge V_{j}\right) \cdot \operatorname{Pr}\left(\neg v_{i} \mid V_{j}\right)
$$

Now consider once more Figure 4.2 showing a fragment of the directed tree $G$ of the network. Using the d-separation criterion, we observe that $\langle\boldsymbol{X}|\left\{V_{i}\right\}\left|\left\{V_{j}\right\}\right\rangle{ }_{G}^{d}$ for all sets of vertices $\boldsymbol{X} \subseteq \boldsymbol{V}_{\boldsymbol{i}}^{-}$. Since $G$ is a directed I-map for the distribution Pr, it follows that $I_{\operatorname{Pr}}\left(\boldsymbol{X},\left\{V_{i}\right\},\left\{V_{j}\right\}\right)$ for all sets $\boldsymbol{X} \subseteq \boldsymbol{V}_{\boldsymbol{i}}^{-}$. So,

$$
\lambda_{V_{i}}^{V_{j}}\left(V_{j}\right)=\operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{\boldsymbol{i}}^{-}} \mid v_{i}\right) \cdot \operatorname{Pr}\left(v_{i} \mid V_{j}\right)+\operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{\boldsymbol{i}}^{-}} \mid \neg v_{i}\right) \cdot \operatorname{Pr}\left(\neg v_{i} \mid V_{j}\right)
$$

The probabilities $\operatorname{Pr}\left(v_{i} \mid V_{j}\right)$ and $\operatorname{Pr}\left(\neg v_{i} \mid V_{j}\right)$ have been specified as function values of the probability assessment function $\gamma_{V_{i}}$ associated with vertex $V_{i}$ and hence are directly available to $V_{i}$. In addition, $V_{i}$ computes the probabilities $\operatorname{Pr}\left(\widetilde{c}_{V_{i}} \mid v_{i}\right)$ and $\operatorname{Pr}\left(\widetilde{c}_{V_{i}^{-}} \mid \neg v_{i}\right)$ as function values of its compound diagnostic parameter $\lambda_{V_{i}}$. The property stated in the lemma now follows by substitution.

Similarly, a vertex can compute the causal parameter to send to a successor from its compound causal parameter and the diagnostic parameters it receives from its other successors. The following lemma states this property more formally.

Lemma 4.2.7 Let $\mathcal{B}=(G, \Gamma)$ be a Bayesian network where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is a directed tree. Let $V_{i} \in \boldsymbol{V}_{G}$ be an uninstantiated vertex with the successors $\boldsymbol{\sigma}_{G}\left(V_{i}\right)=$ $\left\{V_{i_{1}}, \ldots, V_{i_{m}}\right\}, m \geq 1$, in $G$. Furthermore, let $\pi_{V_{i}}$ be the compound causal parameter for vertex $V_{i}$; for each $V_{i_{j}} \in \sigma_{G}\left(V_{i}\right)$, let $\pi_{V_{i_{j}}}^{V_{i}}$ be the causal parameter from $V_{i}$ to $V_{i_{j}}$ and let $\lambda_{V_{i_{j}}}^{V_{i}}$ be the diagnostic parameter from $V_{i_{j}}$ to $V_{i}$. Then,

$$
\pi_{V_{i_{j}}}^{V_{i}}\left(V_{i}\right)=\alpha \cdot \pi_{V_{i}}\left(V_{i}\right) \cdot \prod_{k=1, \ldots, m, k \neq j} \lambda_{V_{i_{k}}}^{V_{i}}\left(V_{i}\right)
$$

where $\alpha$ is a normalisation constant.
Proof. Let Pr be the joint probability distribution defined by the Bayesian network $\mathcal{B}$. For each vertex $V \in \boldsymbol{V}_{G}$, let $\boldsymbol{V}^{+}$and $\boldsymbol{V}^{-}$be as before. Then, by definition we have that

$$
\pi_{V_{i_{j}}}^{V_{i}}\left(V_{i}\right)=\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{i_{j}}^{+}}\right)
$$

Using Bayes' Theorem, we find

$$
\pi_{V_{i_{j}}}^{V_{i}}\left(V_{i}\right)=\frac{\operatorname{Pr}\left(\widetilde{c}_{V_{i_{j}}^{+}} \mid V_{i}\right) \cdot \operatorname{Pr}\left(V_{i}\right)}{\operatorname{Pr}\left(\widetilde{c}_{V_{i_{j}}^{+}}\right)}
$$

Now consider Figure 4.4 showing a fragment of the directed tree $G$ of the network. Since $V_{i}$ is an uninstantiated vertex, we have that $\widetilde{c}_{V_{i_{j}}}=\widetilde{c}_{V_{i}^{+}} \wedge\left(\bigwedge_{k=1, \ldots, m, k \neq j} \widetilde{c}_{V_{i_{k}}}\right)$. So,

$$
\pi_{V_{i_{j}}}^{V_{i}}\left(V_{i}\right)=\frac{\operatorname{Pr}\left(\widetilde{c}_{V_{i}^{+}} \wedge\left(\bigwedge_{k=1, \ldots, m, k \neq j} \widetilde{c}_{V_{i_{k}}^{-}}\right) \mid V_{i}\right) \cdot \operatorname{Pr}\left(V_{i}\right)}{\operatorname{Pr}\left(\widetilde{c}_{V_{i_{j}}}^{+}\right)}
$$

Furthermore, using the d-separation criterion, we observe that $\langle\boldsymbol{X}|\left\{V_{i}\right\}|\boldsymbol{Y}\rangle_{G}^{d}$ for all sets of vertices $\boldsymbol{X} \subseteq \boldsymbol{V}_{\boldsymbol{i}}^{+}$and $\boldsymbol{Y} \subseteq \boldsymbol{V}_{\boldsymbol{i}_{k}}^{-}, k=1, \ldots, m$, and for all sets of vertices $\boldsymbol{X} \subseteq \boldsymbol{V}_{\boldsymbol{i}_{\boldsymbol{k}}}^{-}$and $\boldsymbol{Y} \subseteq \boldsymbol{V}_{\boldsymbol{i}_{l}}^{-}, k=1, \ldots, m, l=1, \ldots, m, k \neq l$. Exploiting this observation, we find

$$
\begin{aligned}
\pi_{V_{i_{j}}}^{V_{i}}\left(V_{i}\right) & =\frac{\operatorname{Pr}\left(\widetilde{c}_{V_{i}^{+}} \mid V_{i}\right) \cdot \prod_{k=1, \ldots, m, k \neq j} \operatorname{Pr}\left(\widetilde{c}_{V_{i_{\boldsymbol{k}}}} \mid V_{i}\right) \cdot \operatorname{Pr}\left(V_{i}\right)}{\operatorname{Pr}\left(\widetilde{c}_{V_{i_{j}}^{+}}\right)}= \\
& =\frac{\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{i}^{+}}\right) \cdot \prod_{k=1, \ldots, m, k \neq j} \operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{\boldsymbol{i}_{\boldsymbol{k}}^{-}}} \mid V_{i}\right) \cdot \operatorname{Pr}\left(\widetilde{c}_{V_{i}^{+}}\right)}{\operatorname{Pr}\left(\widetilde{c}_{V_{\boldsymbol{i}_{j}}^{+}}\right)}
\end{aligned}
$$



Figure 4.4: Exploiting d-Separation for Computing $\pi_{V_{i_{j}}}^{V_{i}}\left(V_{i}\right)$.

The probabilities $\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{i}^{+}}\right)$equal the function values of the compound causal parameter $\pi_{V_{i}}$ for $V_{i}$. The probabilities $\operatorname{Pr}\left(\widetilde{c}_{V_{i_{k}}} \mid V_{i}\right)$ equal the function values of the diagnostic parameter $\lambda_{V_{i_{k}}}^{V_{i}}$ vertex $V_{i}$ receives from its successor $V_{i_{k}}$. In addition, we observe that the factor

$$
\frac{\operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{i}^{+}}\right)}{\operatorname{Pr}\left(\tilde{c}_{\boldsymbol{V}_{i_{j}}^{+}}\right)}=\frac{1}{\operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{i_{j}}^{+}}^{+}{\widetilde{V_{\boldsymbol{V}_{i}}}}^{+}\right.}
$$

is dependent on the location of the variables $V_{i}$ and $V_{i_{j}}$ in the digraph but not on their values. This factor may therefore be looked upon as a normalisation constant for $V_{i}$ and $V_{i_{j}}$, and will henceforth be denoted as $\alpha$; it can be computed from $\operatorname{Pr}\left(v_{i} \mid \widetilde{c}_{V_{i_{j}}}\right)+\operatorname{Pr}\left(\neg v_{i} \mid\right.$ $\left.\widetilde{c}_{V_{i_{j}}^{+}}\right)=1$. The property stated in the lemma now follows by substitution.

The previous lemma applies to uninstantiated vertices only; the lemma, however, can be taken to hold for instantiated vertices as described before.

The data fusion lemma and the four computation rules provided by the Lemmas 4.2.4, 4.2.5, 4.2.6, and 4.2.7 with each other constitute Pearl's algorithm for probabilistic inference with a Bayesian network comprising a directed tree for its qualitative part. Note that Pearl's algorithm provides for computing probabilities as well as for processing evidence in such a Bayesian network. The computation rules for the separate causal and diagnostic parameters enable a vertex to pass on the impact of a piece of evidence correctly, and allow for the evidence to spread throughout the network. Close examination of the computation rules from the Lemmas 4.2.6 and 4.2.7 further reveals that a vertex that sends an updated parameter will not receive a new parameter originating from the same evidence. A causal parameter or a diagnostic parameter to a vertex is not affected by the diagnostic parameter or the causal parameter, respectively, from that vertex. In addition, the topological property that in a directed tree there is at most one chain between any two vertices prohibits the process of parameter updating that originates from some vertex to reach this vertex along another chain.

These properties with each other guarantee that feedback and circular reasoning are prevented and that evidence is propagated throughout the network in a single pass.

We state an additional property concerning the compound diagnostic parameter for a vertex that is useful for investigating the spreading of evidence.

Lemma 4.2.8 Let $\mathcal{B}=(G, \Gamma)$ be a Bayesian network where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is a directed tree. For each vertex $V_{i} \in \boldsymbol{V}_{G}$, let $\lambda_{V_{i}}\left(V_{i}\right)$ be the compound diagnostic parameter for $V_{i}$. If $\widetilde{c}_{V_{G}}=$ True, then $\lambda_{V_{i}}\left(V_{i}\right)=1$ for all $V_{i} \in \boldsymbol{V}_{G}$.

Proof. The property stated in the lemma is proven by (reverse) induction on the depth of the directed tree $G$. Let $n$ be the maximal depth of the tree.

## Induction Basis

The property holds for every leaf of the tree at depth $n$ by definition.

## Induction Hypothesis

For a specific $d \leq n$, we assume that $\lambda_{V_{i}}\left(V_{i}\right)=1$ for all vertices $V_{i}$ at $\operatorname{depth} d, d+$ $1, \ldots, n$.

## Induction Step

Now consider a vertex $V_{i}$ at depth $d-1$ in the tree. We distinguish between two cases. If $V_{i}$ is a leaf of the tree, then $\lambda_{V_{i}}\left(V_{i}\right)=1$ by definition. Now, suppose that $V_{i}$ has $m$ successors $V_{i_{1}}, \ldots, V_{i_{m}}, m \geq 1$. From $\widetilde{c}_{V_{G}}=$ True, it follows that $V_{i}$ is an uninstantiated vertex. Therefore, it follows from the property stated in Lemma 4.2.5 that

$$
\lambda_{V_{i}}\left(V_{i}\right)=\prod_{j=1, \ldots, m} \lambda_{V_{i_{j}}}^{V_{i}}\left(V_{i}\right)
$$

where $\lambda_{V_{i_{j}}}^{V_{i}}$ is the diagnostic parameter from $V_{i_{j}}$ to $V_{i}, j=1, \ldots, m$. For each parameter $\lambda_{V_{i_{j}}}^{V_{i}}$ we have from the property stated in Lemma 4.2 .6 that

$$
\lambda_{V_{i_{j}}}^{V_{i}}\left(V_{i}\right)=\sum_{c_{V_{i_{j}}}} \lambda_{V_{i_{j}}}\left(c_{V_{i_{j}}}\right) \cdot \gamma_{V_{i_{j}}}\left(c_{V_{i_{j}}} \mid V_{i}\right)
$$

Since vertex $V_{i_{j}}$ is a successor of $V_{i}$, it is located at depth $d$ in the directed tree $G$. From the induction hypothesis we have that $\lambda_{V_{i_{j}}}\left(V_{i_{j}}\right)=1$. So,

$$
\lambda_{V_{i_{j}}}^{V_{i}}\left(V_{i}\right)=\sum_{c_{V_{i_{j}}}} \gamma_{V_{i_{j}}}\left(c_{V_{i_{j}}} \mid V_{i}\right)=1
$$

From $\lambda_{V_{i_{j}}}^{V_{i}}\left(V_{i}\right)=1$ for all successors $V_{i_{j}}$ of $V_{i}$, it follows that $\lambda_{V_{i}}\left(V_{i}\right)=1$. Since vertex $V_{i}$ has been chosen arbitrarily, it follows that for each vertex $V_{i} \in \boldsymbol{V}_{G}$, we have $\lambda_{V_{i}}\left(V_{i}\right)=1$.

It will be evident that the previous lemma can be taken to apply to subtrees of a directed tree as well.

We conclude our discussion of Pearl's algorithm for probabilistic inference so far with an example.


Figure 4.5: An Example Bayesian Network.

Example 4.2.9 We consider the Bayesian network $\mathcal{B}=(G, \Gamma)$ shown in Figure 4.5. Let $\operatorname{Pr}$ be the joint probability distribution defined by $\mathcal{B}$. We address the computation of the probabilities and different parameters for the various vertices in the network.

Vertex $V_{1}$ sets out to compute the probabilities of its values by application of the data fusion lemma:

$$
\begin{array}{ll}
\operatorname{Pr}\left(v_{1}\right) & =\alpha \cdot \pi_{V_{1}}\left(v_{1}\right) \cdot \lambda_{V_{1}}\left(v_{1}\right) \\
\operatorname{Pr}\left(\neg v_{1}\right) & =\alpha \cdot \pi_{V_{1}}\left(\neg v_{1}\right) \cdot \lambda_{V_{1}}\left(\neg v_{1}\right)
\end{array}
$$

To be able to compute the probabilities of interest from the lemma, vertex $V_{1}$ needs to compute its compound parameters $\pi_{V_{1}}$ and $\lambda_{V_{1}}$. Since no evidence has been entered into the network as yet, we have from the property stated in Lemma 4.2.8 that the values of the compound diagnostic parameter $\lambda_{V_{1}}$ equal

$$
\begin{array}{ll}
\lambda_{V_{1}}\left(v_{1}\right) & =1 \\
\lambda_{V_{1}}\left(\neg v_{1}\right) & =1
\end{array}
$$

Vertex $V_{1}$ now turns to the computation of its compound causal parameter $\pi_{V_{1}}$. Using the property stated in Lemma 4.2.4, it finds the values of this parameter to be

$$
\begin{array}{ll}
\pi_{V_{1}}\left(v_{1}\right) & =0.5 \\
\pi_{V_{1}}\left(\neg v_{1}\right) & =0.5
\end{array}
$$

Substitution of the values of the compound parameters into the data fusion lemma and subsequent elimination of the normalisation constant $\alpha$ yields

$$
\begin{array}{ll}
\operatorname{Pr}\left(v_{1}\right) & =0.5 \\
\operatorname{Pr}\left(\neg v_{1}\right) & =0.5
\end{array}
$$

Now observe that vertex $V_{1}$ not just computes its probabilities, it also computes the parameter to send to its neighbour in the tree. Using the property stated in Lemma 4.2.7, it computes the values of the causal parameter $\pi_{V_{2}}^{V_{1}}$ for its successor $V_{2}$ to be

$$
\begin{array}{ll}
\pi_{V_{2}}^{V_{1}}\left(v_{1}\right) & =0.5 \\
\pi_{V_{2}}\left(\neg v_{1}\right) & =0.5
\end{array}
$$

By sending the thus computed parameter to vertex $V_{2}$, vertex $V_{1}$ provides $V_{2}$ with sufficient information about the represented probability distribution to enable vertex $V_{2}$ to compute its probabilities.

Just as vertex $V_{1}$, vertex $V_{2}$ sets out to compute the probabilities of its values by application of the data fusion lemma:

$$
\begin{array}{ll}
\operatorname{Pr}\left(v_{2}\right) & =\alpha \cdot \pi_{V_{2}}\left(v_{2}\right) \cdot \lambda_{V_{2}}\left(v_{2}\right) \\
\operatorname{Pr}\left(\neg v_{2}\right) & =\alpha \cdot \pi_{V_{2}}\left(\neg v_{2}\right) \cdot \lambda_{V_{2}}\left(\neg v_{2}\right)
\end{array}
$$

Vertex $V_{2}$ now needs to compute its compound parameters $\pi_{V_{2}}$ and $\lambda_{V_{2}}$. Since no evidence has been entered into the network as yet, we once more have from the property stated in Lemma 4.2.8 that the values of the compound diagnostic parameter equal

$$
\begin{array}{ll}
\lambda_{V_{2}}\left(v_{2}\right) & =1 \\
\lambda_{V_{2}}\left(\neg v_{2}\right) & =1
\end{array}
$$

Using the property stated in Lemma 4.2.4, vertex $V_{2}$ now computes the compound causal parameter $\pi_{V_{2}}$ from its own probability assessment function $\gamma_{V_{2}}$ and the causal parameter $\pi_{V_{2}}^{V_{1}}$ it receives from its predecessor. For the values $\pi_{V_{2}}\left(v_{2}\right)$ and $\pi_{V_{2}}\left(\neg v_{2}\right)$, it thus finds

$$
\begin{aligned}
\pi_{V_{2}}\left(v_{2}\right) & =\gamma_{V_{2}}\left(v_{2} \mid v_{1}\right) \cdot \pi_{V_{2}}^{V_{1}}\left(v_{1}\right)+\gamma_{V_{2}}\left(v_{2} \mid \neg v_{1}\right) \cdot \pi_{V_{2}}^{V_{1}}\left(\neg v_{1}\right)= \\
& =0.8 \cdot 0.5+0.4 \cdot 0.5=0.6 \\
\pi_{V_{2}}\left(\neg v_{2}\right) & =\gamma_{V_{2}}\left(\neg v_{2} \mid v_{1}\right) \cdot \pi_{V_{2}}^{V_{1}}\left(v_{1}\right)+\gamma_{V_{2}}\left(\neg v_{2} \mid \neg v_{1}\right) \cdot \pi_{V_{2}}^{V_{1}}\left(\neg v_{1}\right)= \\
& =0.2 \cdot 0.5+0.6 \cdot 0.5=0.4
\end{aligned}
$$

Substitution of the values of the compound parameters into the data fusion lemma and subsequent elimination of the normalisation constant $\alpha$ yields

$$
\begin{array}{ll}
\operatorname{Pr}\left(v_{2}\right) & =0.6 \\
\operatorname{Pr}\left(\neg v_{2}\right) & =0.4
\end{array}
$$

In addition to computing its own probabilities, vertex $V_{2}$ computes the parameters to send to its neighbours in the tree. For its predecessor $V_{1}$, vertex $V_{2}$ computes the diagnostic parameter $\lambda_{V_{2}}^{V_{1}}$; the values of this parameter equal

$$
\begin{array}{ll}
\lambda_{V_{2}}^{V_{1}}\left(v_{1}\right) & =1 \\
\lambda_{V_{2}}^{V_{1}}\left(\neg v_{1}\right) & =1
\end{array}
$$

Vertex $V_{2}$ further computes a causal parameter for every one of its successors in the tree. Using the property stated in Lemma 4.2.7, it computes the values of the causal parameter $\pi_{V_{3}}^{V_{2}}$ for vertex $V_{3}$ from

$$
\begin{array}{ll}
\pi_{V_{3}}^{V_{2}}\left(v_{2}\right) & =\alpha \cdot \pi_{V_{2}}\left(v_{2}\right) \cdot \lambda_{V_{4}}^{V_{2}}\left(v_{2}\right) \\
\pi_{V_{3}}^{V_{2}}\left(\neg v_{2}\right) & =\alpha \cdot \pi_{V_{2}}\left(\neg v_{2}\right) \cdot \lambda_{V_{4}}^{V_{2}}\left(\neg v_{2}\right)
\end{array}
$$

It is readily seen that the values of the diagnostic parameter $\lambda_{V_{4}}^{V_{2}}$ that vertex $V_{2}$ receives from vertex $V_{4}$ both equal one. Substitution of the various parameter values and subsequent elimination of the normalisation constant $\alpha$ now yields

$$
\begin{array}{ll}
\pi_{V_{3}}^{V_{2}}\left(v_{2}\right) & =0.6 \\
\pi_{V_{3}}^{V_{2}}\left(\neg v_{2}\right) & =0.4
\end{array}
$$

For vertex $V_{4}$, vertex $V_{2}$ equally computes

$$
\begin{array}{ll}
\pi_{V_{4}}^{V_{2}}\left(v_{2}\right) & =0.6 \\
\pi_{V_{4}}^{V_{2}}\left(\neg v_{2}\right) & =0.4
\end{array}
$$

By sending the thus computed parameters to its respective neighbours, vertex $V_{2}$ enables these neighbours to compute their probabilities in turn. Vertex $V_{3}$ now computes the probabilities of its values to be

$$
\begin{array}{ll}
\operatorname{Pr}\left(v_{3}\right) & =0.7 \\
\operatorname{Pr}\left(\neg v_{3}\right) & =0.3
\end{array}
$$

Vertex $V_{4}$ computes its probabilities to be

$$
\begin{array}{ll}
\operatorname{Pr}\left(v_{4}\right) & =0.2 \\
\operatorname{Pr}\left(\neg v_{4}\right) & =0.8
\end{array}
$$

Now, suppose that the evidence $V_{4}=$ true is observed and entered into the Bayesian network $\mathcal{B}$. We address the computation of the posterior probabilities and the different parameters for the various vertices in the network using Pearl's algorithm. Vertex $V_{4}$ once more sets out to compute the probabilities of its values by application of the data fusion lemma:

$$
\begin{array}{ll}
\operatorname{Pr}^{v_{4}}\left(v_{4}\right) & =\alpha \cdot \pi_{V_{4}}\left(v_{4}\right) \cdot \lambda_{V_{4}}\left(v_{4}\right) \\
\operatorname{Pr}^{4}\left(\neg v_{4}\right) & =\alpha \cdot \pi_{V_{4}}\left(\neg v_{4}\right) \cdot \lambda_{V_{4}}\left(\neg v_{4}\right)
\end{array}
$$

Since the evidence $V_{4}=$ true has been entered for vertex $V_{4}$, it finds the values of its compound diagnostic parameter $\lambda_{V_{4}}$ to be

$$
\begin{aligned}
& \lambda_{V_{4}}\left(v_{4}\right)=1 \\
& \lambda_{V_{4}}\left(\neg v_{4}\right)=0
\end{aligned}
$$

resulting in

$$
\begin{array}{ll}
\operatorname{Pr}^{v_{4}}\left(v_{4}\right) & =1 \\
\operatorname{Pr}^{v_{4}}\left(\neg v_{4}\right) & =0
\end{array}
$$

as expected. In addition to updating its own probabilities, vertex $V_{4}$ computes a new parameter to send to its neighbour in the tree. Using the property stated in Lemma 4.2.6, vertex $V_{4}$ computes the values of the diagnostic parameter $\lambda_{V_{4}}^{V_{2}}$ for its predecessor $V_{2}$ to be

$$
\begin{aligned}
\lambda_{V_{4}}^{V_{2}}\left(v_{2}\right) & =\gamma_{V_{4}}\left(v_{4} \mid v_{2}\right) \cdot \lambda_{V_{4}}\left(v_{4}\right)+\gamma_{V_{4}}\left(\neg v_{4} \mid v_{2}\right) \cdot \lambda_{V_{4}}\left(\neg v_{4}\right)= \\
& =0.3 \cdot 1+0.7 \cdot 0=0.3 \\
\lambda_{V_{4}}^{V_{2}}\left(\neg v_{2}\right) & =\gamma_{V_{4}}\left(v_{4} \mid \neg v_{2}\right) \cdot \lambda_{V_{4}}\left(v_{4}\right)+\gamma_{V_{4}}\left(\neg v_{4} \mid \neg v_{2}\right) \cdot \lambda_{V_{4}}\left(\neg v_{4}\right)= \\
& =0.05 \cdot 1+0.95 \cdot 0=0.05
\end{aligned}
$$

By sending the thus computed parameter to vertex $V_{2}$, vertex $V_{4}$ informs $V_{2}$ of the newly entered evidence, enabling it to update its probabilities.

As before, vertex $V_{2}$ sets out to compute the probabilities of its values by application of the data fusion lemma:

$$
\begin{aligned}
& \operatorname{Pr}^{v_{4}}\left(v_{2}\right)=\alpha \cdot \pi_{V_{2}}\left(v_{2}\right) \cdot \lambda_{V_{2}}\left(v_{2}\right) \\
& \operatorname{Pr}^{v_{4}}\left(\neg v_{2}\right)=\alpha \cdot \pi_{V_{2}}\left(\neg v_{2}\right) \cdot \lambda_{V_{2}}\left(\neg v_{2}\right)
\end{aligned}
$$

Since no evidence has been entered for vertex $V_{2}$ 's non-descendants, it once more computes the values of its compound causal parameter $\pi_{V_{2}}$ to be

$$
\begin{array}{ll}
\pi_{V_{2}}\left(v_{2}\right) & =0.6 \\
\pi_{V_{2}}\left(\neg v_{2}\right) & =0.4
\end{array}
$$

Using the property stated in Lemma 4.2.5, vertex $V_{2}$ computes the values of its compound diagnostic parameter to be

$$
\begin{array}{ll}
\lambda_{V_{2}}\left(v_{2}\right) & =\lambda_{V_{3}}^{V_{2}}\left(v_{2}\right) \cdot \lambda_{V_{4}}^{V_{2}}\left(v_{2}\right) \\
\lambda_{V_{2}}\left(\neg v_{2}\right) & =\lambda_{V_{3}}^{V_{2}}\left(\neg v_{2}\right) \cdot \lambda_{V_{4}}^{V_{2}}\left(\neg v_{2}\right)
\end{array}
$$

It is readily seen that the values of the diagnostic parameter $\lambda_{V_{3}}^{V_{2}}$ that vertex $V_{2}$ receives from vertex $V_{3}$ both equal one. Substitution of the various parameter values now yields

$$
\begin{array}{ll}
\lambda_{V_{2}}\left(v_{2}\right) & =0.3 \\
\lambda_{V_{2}}\left(\neg v_{2}\right) & =0.05
\end{array}
$$

Vertex $V_{2}$ substitutes the values of its compound parameters $\pi_{V_{2}}$ and $\lambda_{V_{2}}$ into the data fusion lemma to find

$$
\begin{array}{ll}
\operatorname{Pr}^{v_{4}}\left(v_{2}\right) & =\alpha \cdot 0.6 \cdot 0.3=\alpha \cdot 0.18 \\
\operatorname{Pr}^{v_{4}}\left(\neg v_{2}\right) & =\alpha \cdot 0.4 \cdot 0.05=\alpha \cdot 0.02
\end{array}
$$

After eliminating the normalisation constant $\alpha$ it finds

$$
\begin{array}{ll}
\operatorname{Pr}^{v_{4}}\left(v_{2}\right) & =0.9 \\
\operatorname{Pr}^{v_{4}}\left(\neg v_{2}\right) & =0.1
\end{array}
$$

In addition to updating its own probabilities, vertex $V_{2}$ computes new parameters to send to its other neighbours than vertex $V_{4}$. Using the property stated in Lemma 4.2.7, vertex $V_{2}$ computes the values of the causal parameter $\pi_{V_{3}}^{V_{2}}$ for its successor $V_{3}$ to be

$$
\begin{aligned}
& \pi_{V_{3}}^{V_{2}}\left(v_{2}\right)=\alpha \cdot \pi_{V_{2}}\left(v_{2}\right) \cdot \lambda_{V_{4}}^{V_{2}}\left(v_{2}\right)=\alpha \cdot 0.6 \cdot 0.3=0.9 \\
& \pi_{V_{3}}^{2}\left(\neg v_{2}\right)=\alpha \cdot \pi_{V_{2}}\left(\neg v_{2}\right) \cdot \lambda_{V_{4}}^{V_{2}}\left(\neg v_{2}\right)=\alpha \cdot 0.4 \cdot 0.05=0.1
\end{aligned}
$$

Using the property stated in Lemma 4.2.6, vertex $V_{2}$ computes the values of the diagnostic parameter $\lambda_{V_{2}}^{V_{1}}$ for its predecessor $V_{1}$ to be

$$
\begin{aligned}
\lambda_{V_{2}}^{V_{1}}\left(v_{1}\right) & =\gamma_{V_{2}}\left(v_{2} \mid v_{1}\right) \cdot \lambda_{V_{2}}\left(v_{2}\right)+\gamma_{V_{2}}\left(\neg v_{2} \mid v_{1}\right) \cdot \lambda_{V_{2}}\left(\neg v_{2}\right)= \\
& =0.8 \cdot 0.3+0.2 \cdot 0.05=0.25 \\
\lambda_{V_{2}}^{V_{1}}\left(\neg v_{1}\right) & =\gamma_{V_{2}}\left(v_{2} \mid \neg v_{1}\right) \cdot \lambda_{V_{2}}\left(v_{2}\right)+\gamma_{V_{2}}\left(\neg v_{2} \mid \neg v_{1}\right) \cdot \lambda_{V_{2}}\left(\neg v_{2}\right)= \\
& =0.4 \cdot 0.3+0.6 \cdot 0.05=0.15
\end{aligned}
$$

By sending the thus computed parameters to vertex $V_{3}$ and vertex $V_{1}$, respectively, vertex $V_{2}$ enables these vertices to compute their updated probabilities. Vertex $V_{3}$ computes the updated probabilities of its values to be

$$
\begin{array}{ll}
\operatorname{Pr}^{v_{4}}\left(v_{3}\right) & =0.85 \\
\operatorname{Pr}^{v_{4}}\left(\neg v_{3}\right) & =0.15
\end{array}
$$

Vertex $V_{1}$ computes its updated probabilities to be

$$
\begin{array}{ll}
\operatorname{Pr}^{v_{4}}\left(v_{1}\right) & =0.625 \\
\operatorname{Pr}^{v_{4}}\left(\neg v_{1}\right) & =0.375
\end{array}
$$

### 4.2.2 Singly Connected Digraphs

So far, we have only discussed Pearl's algorithm for probabilistic inference for Bayesian networks comprising a directed tree for their qualitative part. In this section, we extend the algorithm to apply to Bayesian networks of which the qualitative part is a singly connected digraph more in general. Before detailing the computations involved in the extended algorithm, we introduce some new terminology and notational conventions that we will use in the sequel. We consider a singly connected digraph $G$. For this graph $G$, we observe that removal of any arc splits the graph into two separate components. From this property we have that in the digraph $G$ we can identify for a vertex $V_{i}$ with $m$ neighbours, $m$ subgraphs of $G$ each containing a neighbour of $V_{i}$ such that, after removal of $V_{i}$ and all its incoming and outgoing arcs, there does not exist a path from one such subgraph to another. The following definition introduces these subgraphs more formally; Figure 4.6 illustrates the basic idea.


Figure 4.6: Upper and Lower Graphs.

Definition 4.2.10 Let $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ be a singly connected directed graph. For each $\operatorname{arc}\left(V_{i}, V_{j}\right) \in \boldsymbol{A}_{G}$, let $G_{\left(V_{i}, V_{j}\right)}=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G} \backslash\left\{\left(V_{i}, V_{j}\right)\right\}\right)$ be the digraph that results after removing arc $\left(V_{i}, V_{j}\right)$ from $G$. Now, let $V_{i} \in \boldsymbol{V}_{G}$ be a vertex in $G$. For each predecessor $V_{j} \in \rho_{G}\left(V_{i}\right)$ of $V_{i}$, let $G_{\left(V_{j}, V_{i}\right)}^{+}$be the component of $G_{\left(V_{j}, V_{i}\right)}$ that includes $V_{j}$ in its vertex set; $G_{\left(V_{j}, V_{i}\right)}^{+}$is called an upper graph of vertex $V_{i}$. For each successor $V_{k} \in \sigma_{G}\left(V_{i}\right)$ of $V_{i}$, let $G_{\left(V_{i}, V_{k}\right)}^{-}$be the component of $G_{\left(V_{i}, V_{k}\right)}$ that includes $V_{k}$ in its vertex set; $G_{\left(V_{i}, V_{k}\right)}^{-}$ is called a lower graph of $V_{i}$.

In the previous section, we have detailed Pearl's computation rules for probabilistic inference with a Bayesian network comprising a directed tree: we recall that these rules address the computation of the different parameters for the various vertices in the tree. The proofs of the lemmas that we have presented show that these computation rules derive from exploiting independences. These independences are read from the
qualitative part of a Bayesian network by local inspection of a vertex' incoming and outgoing arcs only. The computation rules therefore make use explicitly of the property that in the network's digraph there is at most one chain between any two vertices. Since singly connected digraphs share this property with directed trees, Pearl's tree algorithm is extended straightforwardly to apply to Bayesian networks comprising a singly connected digraph for their qualitative part. In fact, only the computation rules for the compound causal parameter and for the separate diagnostic parameters are adapted to account for a vertex having multiple predecessors; all other computation rules remain unaltered. For the sake of completeness, we will review all computation rules involved.

We begin by addressing the computation of probabilities from a Bayesian network comprising a singly connected digraph. It will be evident that the probabilities of the values of a vertex of interest are dependent upon all evidence entered into the network so far. As we have seen for the vertices in a directed tree, the probabilities of the values of a vertex in a singly connected digraph can be written in terms of two factors. These factors describe the combined influence on the vertex' probabilities of evidence entered for the vertices in its upper graphs and of evidence entered for the vertices in its lower graphs, respectively. We redefine the compound causal and diagnostic parameters for a vertex to capture these factors in view of singly connected digraphs.

Definition 4.2.11 Let $\mathcal{B}=(G, \Gamma)$ be a Bayesian network where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is a singly connected digraph, and let $\operatorname{Pr}$ be the joint probability distribution defined by $\mathcal{B}$. Let $V_{i} \in \boldsymbol{V}_{G}$ be a vertex in $G$, and let $\boldsymbol{V}_{\boldsymbol{i}}^{+}=\bigcup_{V_{j} \in \boldsymbol{\rho}_{G}\left(V_{i}\right)} \boldsymbol{V}_{G_{\left(V_{j}, V_{i}\right)}^{+}}$and $\boldsymbol{V}_{\boldsymbol{i}}^{-}=\boldsymbol{V}_{G} \backslash \boldsymbol{V}_{\boldsymbol{i}}^{+}$. The compound causal parameter $\pi_{V_{i}}$ for vertex $V_{i}$ is the function $\pi_{V_{i}}:\left\{V_{i}\right\} \rightarrow[0,1]$ defined by

$$
\pi_{V_{i}}\left(V_{i}\right)=\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{\boldsymbol{V}_{i}^{+}}\right)
$$

The compound diagnostic parameter $\lambda_{V_{i}}$ for $V_{i}$ is the function $\lambda_{V_{i}}:\left\{V_{i}\right\} \rightarrow[0,1]$ defined by

$$
\lambda_{V_{i}}\left(V_{i}\right)=\operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{\boldsymbol{i}}^{-}} \mid V_{i}\right)
$$

Note that this redefinition of the compound parameters differs from Definition 4.2.2 with respect to the sets $\boldsymbol{V}_{\boldsymbol{i}}^{+}$and $\boldsymbol{V}_{\boldsymbol{i}}^{-}$for a vertex $V_{i}$. The basic idea of the parameters, however, remains unaltered. The compound causal parameter for a vertex describes the combined influence on its probabilities of all evidence that has been entered 'above' it in the digraph; the compound diagnostic parameter for the vertex describes the combined influence of all evidence that has been entered 'below' it in the digraph. Also note that, for a vertex $W$ without any incoming arcs, we once more find $\pi_{W}(W)=\operatorname{Pr}(W)$; for a vertex $V$ without any outgoing arcs, we again have $\lambda_{V}(V)=1$.

Using the redefinition of the compound causal and diagnostic parameters for a vertex, the data fusion lemma now applies to a Bayesian network comprising a singly connected digraph for its qualitative part: for each vertex $V_{i}$ in the network, we once more have

$$
\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{G}}\right)=\alpha \cdot \pi_{V_{i}}\left(V_{i}\right) \cdot \lambda_{V_{i}}\left(V_{i}\right)
$$

where $\alpha$ is a normalisation constant.

The compound diagnostic and causal parameters for a vertex specify probabilistic information from its lower graphs combined and from all its upper graphs combined, respectively. We observe that to be able to exploit the qualitative part of a Bayesian network as a computational architecture, these compound parameters once again have to be computed from separate causal and diagnostic parameters originating from the various neighbours of the vertex. We redefine these separate parameters before addressing the computation of the compound ones.

Definition 4.2.12 Let $\mathcal{B}=(G, \Gamma)$ be a Bayesian network where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is a singly connected digraph, and let $\operatorname{Pr}$ be the joint probability distribution defined by $\mathcal{B}$. Let $V_{i}$ be a vertex with a successor $V_{k}$ in $G$ and let $G_{\left(V_{i}, V_{k}\right)}^{+}$be as before. The causal parameter $\pi_{V_{k}}^{V_{i}}$ from $V_{i}$ to $V_{k}$ is the function $\pi_{V_{k}}^{V_{i}}:\left\{V_{i}\right\} \rightarrow[0,1]$ defined by

$$
\pi_{V_{k}}^{V_{i}}\left(V_{i}\right)=\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{\left.V_{G_{\left(V_{i}, V_{k}\right)}^{+}}\right)}\right)
$$

Now, let $V_{i}$ be a vertex having a predecessor $V_{j}$ in $G$ and let $G_{\left(V_{j}, V_{i}\right)}^{-}$be as before. The diagnostic parameter $\lambda_{V_{i}}^{V_{j}}$ from $V_{i}$ to $V_{j}$ is the function $\lambda_{V_{i}}^{V_{j}}:\left\{V_{j}\right\} \rightarrow[0,1]$ defined by

$$
\lambda_{V_{i}}^{V_{j}}\left(V_{j}\right)=\operatorname{Pr}\left(\widetilde{c}_{V_{G_{\left(V_{j}, V_{i}\right)}}} \mid V_{j}\right)
$$

The separate causal and diagnostic parameters defined above once again are the messages the vertices send one another through the communication channels of the computational architecture and, hence, may again be looked upon as associated with the arcs of the digraph of the Bayesian network at hand.

The following lemma now shows that a vertex can compute its compound causal parameter from its own probability assessment function and the causal parameters it receives from its various predecessors.

Lemma 4.2.13 Let $\mathcal{B}=(G, \Gamma)$ be a Bayesian network where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is a singly connected digraph. Let $V_{i} \in \boldsymbol{V}_{G}$ be a vertex with the predecessors $\boldsymbol{\rho}_{G}\left(V_{i}\right)=$ $\left\{V_{j_{1}}, \ldots, V_{j_{n}}\right\}, n \geq 1$, in $G$. Furthermore, let $\pi_{V_{i}}$ be the compound causal parameter for vertex $V_{i}$ and, for each $V_{j_{k}} \in \boldsymbol{\rho}_{G}\left(V_{i}\right)$, let $\pi_{V_{i}}^{V_{j_{k}}}$ be the causal parameter from $V_{j_{k}}$ to $V_{i}$. Then,

$$
\pi_{V_{i}}\left(V_{i}\right)=\sum_{c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}} \gamma_{V_{i}}\left(V_{i} \mid c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right) \cdot \prod_{k=1, \ldots, n} \pi_{V_{i}}^{V_{j_{k}}}\left(c_{V_{j_{k}}}\right)
$$

where $c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}=\bigwedge_{k=1, \ldots, n} c_{V_{j_{k}}}$.
Proof. Let $\operatorname{Pr}$ be the joint probability distribution defined by the Bayesian network $\mathcal{B}$. For vertex $V_{i}$, let $\boldsymbol{V}_{\boldsymbol{i}}^{+}=\bigcup_{k=1, \ldots, n} \boldsymbol{V}_{G_{\left(V_{j_{k}}, V_{i}\right)}^{+}}$be as before. Then, by definition we have that

$$
\pi_{V_{i}}\left(V_{i}\right)=\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{i}^{+}}\right)=\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{G_{\left(V_{j_{1}}, V_{i}\right)}^{+}}} \wedge \cdots \wedge \widetilde{c}_{V_{\left.G_{\left(V_{j_{n}}, V_{i}\right)}\right)}}\right)
$$



Figure 4.7: Exploiting d-Separation for Computing $\pi_{V_{i}}\left(V_{i}\right)$.

By conditioning on the various configurations of the set $\boldsymbol{\rho}_{G}\left(V_{i}\right)$ of $V_{i}$ 's predecessors, we find

$$
\begin{gathered}
\pi_{V_{i}}\left(V_{i}\right)=\sum_{c_{\rho_{G}\left(V_{i}\right)}} \operatorname{Pr}\left(V_{i} \mid c_{\rho_{G}\left(V_{i}\right)} \wedge \widetilde{c}_{V_{G_{\left(V_{j_{1}}, V_{i}\right)}}} \wedge \cdots \wedge \widetilde{c}_{V_{G_{\left(V_{j_{j}}, V_{i}\right)}}}\right) . \\
\cdot \operatorname{Pr}\left(c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)} \mid \widetilde{c}_{V_{G_{\left(V_{j_{1}}, V_{i}\right)}^{+}}} \wedge \cdots \wedge \widetilde{c}_{V_{G_{\left(V_{j_{n}}, V_{i}\right)}}}\right)
\end{gathered}
$$

Now consider Figure 4.7 showing a fragment of the singly connected digraph $G$ of the network. Using the d-separation criterion, we observe that $\left\langle\left\{V_{i}\right\}\right| \boldsymbol{\rho}_{G}\left(V_{i}\right)|\boldsymbol{X}\rangle_{G}^{d}$ for all sets of vertices $\boldsymbol{X} \subseteq \boldsymbol{V}_{\boldsymbol{i}}^{+}$. Since $G$ is a directed I-map for the distribution Pr, it follows that $I_{\operatorname{Pr}}\left(\left\{V_{i}\right\}, \boldsymbol{\rho}_{G}\left(V_{i}\right), \boldsymbol{X}\right)$ for all sets $\boldsymbol{X} \subseteq \boldsymbol{V}_{\boldsymbol{i}}^{+}$. Exploiting this observation, we find

$$
\pi_{V_{i}}\left(V_{i}\right)=\sum_{c_{\rho_{G}\left(V_{i}\right)}} \operatorname{Pr}\left(V_{i} \mid c_{\rho_{G}\left(V_{i}\right)}\right) \cdot \operatorname{Pr}\left(c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)} \mid \widetilde{c}_{V_{G_{\left(V_{j_{1}}, V_{i}\right)}}} \wedge \cdots \wedge \widetilde{c}_{V_{\left.G_{\left(V_{j},\right.}, V_{i}\right)}}\right)
$$

By once more using the d-separation criterion, we find that $I_{\mathrm{Pr}}\left(\boldsymbol{X},\left\{V_{j_{k}}\right\}, \boldsymbol{Y}\right)$ for all sets of vertices $\boldsymbol{X} \subseteq \boldsymbol{V}_{G_{\left(V_{j_{k}}, V_{i}\right)}^{+}}, \boldsymbol{Y} \subseteq \boldsymbol{V}_{G_{\left(V_{j_{l}}, V_{i}\right)}^{+}}, l=1, \ldots, n, k=1, \ldots, n, k \neq l$. Exploiting this observation, we find

$$
\pi_{V_{i}}\left(V_{i}\right)=\sum_{c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}} \operatorname{Pr}\left(V_{i} \mid c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right) \cdot \operatorname{Pr}\left(c_{V_{j_{1}}} \mid \widetilde{c}_{V_{G_{\left(V_{j_{1}}, V_{i}\right)}}}\right) \cdot \ldots \cdot \operatorname{Pr}\left(c_{V_{j_{n}}} \mid \widetilde{c}_{V_{G_{\left(V_{j_{n}}, V_{i}\right)}^{+}}}\right)
$$

where $c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}=\bigwedge_{k=1, \ldots, n} c_{V_{j_{k}}}$. The probabilities $\operatorname{Pr}\left(V_{i} \mid c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)$ have been specified as function values of the probability assessment function $\gamma_{V_{i}}$ and therefore are directly available to vertex $V_{i}$. In addition, $V_{i}$ receives the probabilities $\operatorname{Pr}\left(c_{V_{j_{k}}} \mid \widetilde{c}_{V_{\left.G_{\left(V_{j_{k}}, V_{i}\right)}\right)}}\right)$
from its predecessor $V_{j_{k}}$ as function values of the causal parameter $\pi_{V_{i}}^{V_{j_{k}}}, k=1, \ldots, n$. The property stated in the lemma now follows by substitution.

A vertex can further compute its compound diagnostic parameter from the separate diagnostic parameters it receives from its successors in the digraph. This property is analogous to the property stated in Lemma 4.2.5, that is, for an uninstantiated vertex $V_{i} \in \boldsymbol{V}_{G}$ with successors $\boldsymbol{\sigma}_{G}\left(V_{i}\right)=\left\{V_{i_{1}}, \ldots, V_{i_{m}}\right\}, m \geq 1$, in the digraph, we have that

$$
\lambda_{V_{i}}\left(V_{i}\right)=\prod_{j=1, \ldots, m} \lambda_{V_{i_{j}}}^{V_{i}}\left(V_{i}\right)
$$

The property can be taken to apply to instantiated vertices as outlined before.
A vertex in turn has to compute parameters to send to its various neighbours. A vertex $V_{i}$ can compute the diagnostic parameter to send to a predecessor $V_{j_{k}}$ from its own probability assessment function, its compound diagnostic parameter, and the causal parameters it receives from its other predecessors.

Lemma 4.2.14 Let $\mathcal{B}=(G, \Gamma)$ be a Bayesian network where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is a singly connected graph. Let $V_{i} \in \boldsymbol{V}_{G}$ be a vertex with predecessors $\boldsymbol{\rho}_{G}\left(V_{i}\right)=\left\{V_{j_{1}}, \ldots, V_{j_{n}}\right\}$, $n \geq 1$, in $G$. Let $\lambda_{V_{i}}$ be the compound diagnostic parameter for vertex $V_{i}$, and let $\pi_{V_{i}}^{V_{j_{l}}}$ be the causal parameter from $V_{j_{l}} \in \rho_{G}\left(V_{i}\right)$ to $V_{i}$. Futhermore, for each $V_{j_{k}} \in \rho_{G}\left(V_{i}\right)$, let $\lambda_{V_{i}}^{V_{j_{k}}}$ be the diagnostic parameter from $V_{i}$ to $V_{j_{k}}$. Then,

$$
\begin{array}{r}
\lambda_{V_{i}}^{V_{j_{k}}}\left(V_{j_{k}}\right)=\alpha \cdot \sum_{c_{V_{i}}} \lambda_{V_{i}}\left(c_{V_{i}}\right) \cdot\left[\sum _ { c _ { \rho _ { G } ( V _ { i } ) \backslash \{ V _ { j _ { k } } \} } } \left(\gamma_{V_{i}}\left(c_{V_{i}} \mid c_{\rho_{G}\left(V_{i}\right) \backslash\left\{V_{j_{k}}\right\}} \wedge V_{j_{k}}\right) .\right.\right. \\
\left.\left.\cdot \prod_{l=1, \ldots, n, l \neq k} \pi_{V_{i}}^{V_{j_{l}}}\left(c_{V_{j_{l}}}\right)\right)\right]
\end{array}
$$

where $c_{\boldsymbol{\rho}_{G}\left(V_{i}\right) \backslash\left\{V_{j_{k}}\right\}}=\bigwedge_{l=1, \ldots, n, l \neq k} c_{V_{j_{l}}}$ and $\alpha$ is a normalisation constant.
Proof. Let Pr be the joint probability distribution defined by the Bayesian network $\mathcal{B}$. Let $\boldsymbol{V}_{\boldsymbol{i}}^{-}, G^{-}$and $G^{+}$be as before. Then, by definition we have that

$$
\left.\lambda_{V_{i}}^{V_{j_{k}}}\left(V_{j_{k}}\right)=\operatorname{Pr}\left(\widetilde{c}_{V_{G_{\left(V_{j_{k}}, V_{i}\right)}}} \mid V_{j_{k}}\right)=\operatorname{Pr}\left(\widetilde{c}_{V_{i}^{-}} \wedge \widetilde{c}_{V_{i}^{+\backslash k}}\right) \mid V_{j_{k}}\right),
$$

where $\boldsymbol{V}_{i}^{+\backslash k}=\bigcup_{h=1, \ldots, n, h \neq k} \boldsymbol{V}_{G_{\left(V_{j_{h}}, V_{i}\right)}^{+}}$captures the vertices in the upper graphs of all predecessors of $V_{i}$ except $V_{j_{k}}$.

By conditioning on the values of $V_{i}$, we now find

$$
\lambda_{V_{i}}^{V_{j_{k}}}\left(V_{j_{k}}\right)=\sum_{c_{V_{i}}} \operatorname{Pr}\left(\widetilde{c}_{V_{i}^{-}} \wedge \widetilde{c}_{\boldsymbol{v}_{i}+\backslash k} \mid c_{V_{i}} \wedge V_{j_{k}}\right) \cdot \operatorname{Pr}\left(c_{V_{i}} \mid V_{j_{k}}\right)
$$

Using d-separation, it now follows that given $V_{i}, \boldsymbol{V}_{\boldsymbol{i}}^{-}$is independent of $\boldsymbol{V} \backslash \boldsymbol{V}_{\boldsymbol{i}}^{-}$. So,

$$
\begin{aligned}
\lambda_{V_{i}}^{V_{j_{k}}}\left(V_{j_{k}}\right) & =\sum_{c_{V_{i}}} \operatorname{Pr}\left(\widetilde{c}_{V_{i}^{-}} \mid c_{V_{i}} \wedge V_{j_{k}}\right) \cdot \operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{i}}+\boldsymbol{k} \mid c_{V_{i}} \wedge V_{j_{k}}\right) \cdot \operatorname{Pr}\left(c_{V_{i}} \mid V_{j_{k}}\right)= \\
& =\sum_{c_{V_{i}}} \operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{i}^{-}} \mid c_{V_{i}}\right) \cdot \operatorname{Pr}\left(c_{V_{i}} \mid V_{j_{k}}\right) \cdot \frac{\operatorname{Pr}\left(c_{V_{i}} \mid \widetilde{c}_{\boldsymbol{V}_{i}} \backslash \boldsymbol{k} \wedge V_{j_{k}}\right) \cdot \operatorname{Pr}\left(\widetilde{c}_{\boldsymbol{V}_{i}+\boldsymbol{k}} \mid V_{j_{k}}\right)}{\operatorname{Pr}\left(c_{V_{i}} \mid V_{j_{k}}\right)}
\end{aligned}
$$

Vertex $V_{i}$ computes the probabilities $\operatorname{Pr}\left(\widetilde{c}_{V_{i}^{-}} \mid c_{V_{i}}\right)$ as function values of its compound diagnostic parameter $\lambda_{V_{i}}$. d-Separation tells us that all predecessors of $V_{i}$ are independent of each other, so

$$
\operatorname{Pr}\left(\widetilde{c}_{V_{i}+\boldsymbol{k}} \mid V_{j_{k}}\right)=\prod_{h=1, \ldots, n, h \neq k} \operatorname{Pr}\left(\widetilde{c}_{V_{G_{\left(V_{j_{h}}, V_{i}\right)}}}\right)
$$

The latter factor is constant for $V_{i}$ and $V_{i_{j}}$ and may therefore be looked upon as a normalisation constant, which we denote by $\alpha .^{2}$

We now focus on the term $\operatorname{Pr}\left(c_{V_{i}} \mid \widetilde{c}_{\boldsymbol{V}_{i}} \backslash \boldsymbol{k} \wedge V_{j_{k}}\right)$ and condition on the values of all predecessors of $V_{i}$ except $V_{j_{k}}$ :

$$
\sum_{c_{\boldsymbol{\rho}_{G}\left(V_{i}\right) \backslash\left\{V_{j_{k}}\right\}}} \operatorname{Pr}\left(c_{V_{i}} \mid \widetilde{c}_{\boldsymbol{V}_{i}^{+} \backslash \boldsymbol{k}} \wedge V_{j_{k}} \wedge c_{\boldsymbol{\rho}_{G}\left(V_{i}\right) \backslash\left\{V_{j_{k}}\right\}}\right) \cdot \operatorname{Pr}\left(c_{\boldsymbol{\rho}_{G}\left(V_{i}\right) \backslash\left\{V_{j_{k}}\right\}} \mid \widetilde{c}_{\boldsymbol{V}_{i}}+\backslash \boldsymbol{k}, ~ \wedge V_{j_{k}}\right)
$$

Using d-separation, it now follows that given $V_{i}$ 's predecessors, $V_{i}$ is independent of $\boldsymbol{V} \backslash \boldsymbol{V}_{\boldsymbol{i}}^{-}$; in addition, recall that all predecessors of $V_{i}$ are independent of each other. So the above term reduces to

$$
\begin{aligned}
& \sum_{c_{\boldsymbol{\rho}_{G}\left(V_{i}\right) \backslash\left\{V_{j_{k}}\right\}}} \operatorname{Pr}\left(c_{V_{i}} \mid V_{j_{k}} \wedge c_{\boldsymbol{\rho}_{G}\left(V_{i}\right) \backslash\left\{V_{j_{k}}\right\}}\right) \cdot \prod_{l=1, \ldots, n, l \neq k} \operatorname{Pr}\left(c_{V_{j_{l}}} \mid \widetilde{c}_{\boldsymbol{V}_{i}+\boldsymbol{k}} \wedge V_{j_{k}}\right) \\
& =\sum_{c_{\boldsymbol{\rho}_{G}\left(V_{i}\right) \backslash\left\{V_{j_{k}}\right\}}} \operatorname{Pr}\left(c_{V_{i}} \mid V_{j_{k}} \wedge c_{\boldsymbol{\rho}_{G}\left(V_{i}\right) \backslash\left\{V_{j_{k}}\right\}}\right) \cdot \prod_{l=1, \ldots, n, l \neq k} \operatorname{Pr}\left(c_{V_{j_{l}}} \mid \widetilde{c}_{V_{G_{\left(V_{j_{l}}, V_{i}\right)}}}\right)
\end{aligned}
$$

The probabilities $\operatorname{Pr}\left(c_{V_{i}} \mid V_{j_{k}} \wedge c_{\boldsymbol{\rho}_{G}\left(V_{i}\right) \backslash\left\{V_{j_{k}}\right\}}\right)$ have been specified as function values of the probability assessment function $\gamma_{V_{i}}$ associated with vertex $V_{i}$ and hence are directly available to $V_{i}$. In addition, $V_{i}$ receives the probabilities $\operatorname{Pr}\left(c_{V_{j_{l}}} \mid \widetilde{c}_{V_{G_{\left(V_{j}, V_{i}\right)}}}\right)$ from each of its predecessors $\boldsymbol{\rho}_{G}\left(V_{i}\right) \backslash\left\{V_{j_{k}}\right\}$ as function values of their causal parameters $\pi_{V_{i}}^{V_{j_{l}}}$. The property stated in the lemma now follows by substitution.

Finally, if we consider a vertex $V_{i}$ with successors $\boldsymbol{\sigma}_{G}\left(V_{i}\right)=\left\{V_{i_{1}}, \ldots, V_{i_{m}}\right\}$ in the digraph, then the causal parameter vertex $V_{i}$ has to send to a successor, is computed from its compound causal parameter and the diagnostic parameters it receives from its other successors; this property is analogous to the property stated in Lemma 4.2.7:

$$
\pi_{V_{i_{j}}}^{V_{i}}\left(V_{i}\right)=\alpha \cdot \pi_{V_{i}}\left(V_{i}\right) \cdot \prod_{k=1, \ldots, m, k \neq j} \lambda_{V_{i_{k}}}^{V_{i}}\left(V_{i}\right)
$$

where $\alpha$ once more is a normalisation constant.
The data fusion lemma and the four computation rules mentioned above with each other constitute Pearl's algorithm for probabilistic inference with a Bayesian network comprising a singly connected digraph for its qualitative part. Once more, these computation rules provide for computing probabilities as well as for processing evidence in such a Bayesian network. The computation rules for the separate causal and diagnostic parameters enable a vertex to correctly pass on the impact of a piece of evidence to its neighbours. As feedback and circular reasoning once more are excluded, evidence is thus propagated throughout the network in a single pass.

Example 4.2.15 We consider once more the Bayesian network $\mathcal{B}=(G, \Gamma)$ from Example 4.1.2; for ease of reference, the network is reproduced in Figure 4.8. Note that the digraph $G$ of $\mathcal{B}$ is not a directed tree, yet is singly connected.

[^4]

Figure 4.8: The Example Bayesian Network Reproduced.
Let $\operatorname{Pr}$ be the joint probability distribution defined by the network $\mathcal{B}$. We address the computation of the probabilities $\operatorname{Pr}\left(v_{3}\right)$ and $\operatorname{Pr}\left(\neg v_{3}\right)$ using Pearl's algorithm. Vertex $V_{3}$ sets out to compute the probabilities of interest by application of the data fusion lemma:

$$
\begin{array}{ll}
\operatorname{Pr}\left(v_{3}\right) & =\alpha \cdot \pi_{V_{3}}\left(v_{3}\right) \cdot \lambda_{V_{3}}\left(v_{3}\right) \\
\operatorname{Pr}\left(\neg v_{3}\right) & =\alpha \cdot \pi_{V_{3}}\left(\neg v_{3}\right) \cdot \lambda_{V_{3}}\left(\neg v_{3}\right)
\end{array}
$$

Vertex $V_{3}$ now needs to compute its compound parameters $\pi_{V_{3}}$ and $\lambda_{V_{3}}$. Since no evidence has been entered into the network as yet, we have that the values of the compound diagnostic parameter $\lambda_{V_{3}}$ equal

$$
\begin{array}{ll}
\lambda_{V_{3}}\left(v_{3}\right) & =1 \\
\lambda_{V_{3}}\left(\neg v_{3}\right) & =1
\end{array}
$$

Vertex $V_{3}$ computes the values of its compound causal parameter $\pi_{V_{3}}$ from

$$
\begin{aligned}
& \pi_{V_{3}}\left(v_{3}\right)= \gamma_{V_{3}}\left(v_{3} \mid v_{1} \wedge v_{2}\right) \cdot \pi_{V_{3}}^{V_{1}}\left(v_{1}\right) \cdot \pi_{V_{3}}^{V_{2}}\left(v_{2}\right)+ \\
&+\gamma_{V_{3}}\left(v_{3} \mid \neg v_{1} \wedge v_{2}\right) \cdot \pi_{V_{3}}^{V_{1}}\left(v_{1}\right) \cdot \pi_{3}^{V_{2}}\left(v_{2}\right)+ \\
&+\gamma_{V_{3}}\left(v_{3} \mid v_{1} \wedge \neg v_{2}\right) \cdot \pi_{V_{3}}^{V_{1}}\left(v_{1}\right) \cdot \pi_{V_{3}}^{V_{3}}\left(\neg v_{2}\right)+ \\
&\left.+\gamma_{V_{3}}\left(v_{3} \mid \neg v_{1} \wedge \neg v_{2}\right) \cdot \pi_{V_{3}}^{V_{1}}\left(\neg v_{1}\right) \cdot \pi_{V_{3}}^{V_{2}} \neg v_{2}\right) \\
&= \gamma_{V_{3}}\left(\neg v_{3} \mid v_{1} \wedge v_{2}\right) \cdot \pi_{V_{3}}^{V_{1}}\left(v_{1}\right) \cdot \pi_{V_{3}}^{V_{2}\left(v_{2}\right)+} \\
&+\gamma_{V_{3}}\left(\neg v_{3} \mid \neg v_{1} \wedge v_{2}\right) \cdot \pi_{V_{3}}^{V_{1}}\left(v_{1}\right) \cdot \pi_{3}^{V_{3}}\left(v_{2}\right)+ \\
& \pi_{V_{3}}\left(\neg v_{3}\right) \\
&+ \\
&+\gamma_{V_{3}}\left(\neg v_{3} \mid v_{1} \wedge \neg v_{2}\right) \cdot \pi_{V_{3}}^{V_{1}}\left(v_{1}\right) \cdot \pi_{V_{3}}^{V_{3}}\left(\neg v_{2}\right)+ \\
& \gamma_{V_{3}}\left(\neg v_{3} \mid \neg v_{1} \wedge \neg v_{2}\right) \cdot \pi_{V_{3}}^{V_{1}}\left(\neg v_{1}\right) \cdot \pi_{V_{3}}^{V_{2}}\left(\neg v_{2}\right)
\end{aligned}
$$

It is readily seen that vertex $V_{3}$ receives from its predecessor $V_{1}$ the causal parameter $\pi_{V_{3}}^{V_{1}}$ with the values

$$
\begin{array}{ll}
\pi_{V_{3}}^{V_{1}}\left(v_{1}\right) & =0.25 \\
\pi_{V_{3}}^{1}\left(\neg v_{1}\right) & =0.75
\end{array}
$$

From its predecessor $V_{2}$, vertex $V_{3}$ receives the causal parameter $\pi_{V_{3}}^{V_{2}}$ with the values

$$
\begin{array}{ll}
\pi_{V_{3}}^{V_{2}}\left(v_{2}\right) & =0.5 \\
\pi_{V_{3}}\left(\neg v_{2}\right) & =0.5
\end{array}
$$

Substitution now yields the following values for $V_{3}$ 's compound causal parameter:

$$
\begin{aligned}
\pi_{V_{3}}\left(v_{3}\right) & =0.75 \cdot 0.25 \cdot 0.5+0.4 \cdot 0.75 \cdot 0.5+0.25 \cdot 0.25 \cdot 0.5+0.2 \cdot 0.75 \cdot 0.5= \\
& =0.35 \\
\pi_{V_{3}}\left(\neg v_{3}\right) & =0.25 \cdot 0.25 \cdot 0.5+0.6 \cdot 0.75 \cdot 0.5+0.75 \cdot 0.25 \cdot 0.5+0.8 \cdot 0.75 \cdot 0.5= \\
& =0.65
\end{aligned}
$$

Subsequent substitution of the values of the compound parameters into the data fusion lemma and elimination of the normalisation constant $\alpha$ yields

$$
\begin{array}{ll}
\operatorname{Pr}\left(v_{3}\right) & =0.35 \\
\operatorname{Pr}\left(\neg v_{3}\right) & =0.65
\end{array}
$$

Now, suppose that the evidence $V_{3}=$ true is observed and entered into the Bayesian network. We address the computation of the posterior probabilities $\operatorname{Pr}^{v_{3}}\left(v_{1}\right)$ and $\operatorname{Pr}^{v_{3}}\left(\neg v_{1}\right)$ using Pearl's algorithm. Vertex $V_{1}$ sets out to compute the probabilities of interest by application of the data fusion lemma:

$$
\begin{array}{ll}
\operatorname{Pr}^{v_{3}}\left(v_{1}\right) & =\alpha \cdot \pi_{V_{1}}\left(v_{1}\right) \cdot \lambda_{V_{1}}\left(v_{1}\right) \\
\operatorname{Pr}^{v_{3}}\left(\neg v_{1}\right) & =\alpha \cdot \pi_{V_{1}}\left(\neg v_{1}\right) \cdot \lambda_{V_{1}}\left(\neg v_{1}\right)
\end{array}
$$

Vertex $V_{1}$ now has to compute its compound parameters $\pi_{V_{1}}$ and $\lambda_{V_{1}}$. For the compound causal parameter, it finds the values

$$
\begin{array}{ll}
\pi_{V_{1}}\left(v_{1}\right) & =0.25 \\
\pi_{V_{1}}\left(\neg v_{1}\right) & =0.75
\end{array}
$$

The values of the compound diagnostic parameter $\lambda_{V_{1}}$ are computed from

$$
\begin{array}{ll}
\lambda_{V_{1}}\left(v_{1}\right) & =\lambda_{V_{3}}^{V_{1}}\left(v_{1}\right) \\
\lambda_{V_{1}}\left(\neg v_{1}\right) & =\lambda_{V_{3}}^{V_{1}}\left(\neg v_{1}\right)
\end{array}
$$

From its successor $V_{3}$, vertex $V_{1}$ receives the diagnostic parameter $\lambda_{V_{3}}^{V_{1}}$ with the values

$$
\begin{aligned}
\lambda_{V_{3}}^{V_{1}}\left(v_{1}\right)= & \lambda_{V_{3}}\left(v_{3}\right) \cdot\left[\gamma_{V_{3}}\left(v_{3} \mid v_{1} \wedge v_{2}\right) \cdot \pi_{V_{3}}^{V_{2}}\left(v_{2}\right)+\gamma_{V_{3}}\left(v_{3} \mid v_{1} \wedge \neg v_{2}\right) \cdot \pi_{V_{3}}^{V_{2}}\left(\neg v_{2}\right)\right] \\
& +\lambda_{V_{3}}\left(\neg v_{3}\right) \cdot\left[\gamma_{V_{3}}\left(\neg v_{3} \mid v_{1} \wedge v_{2}\right) \cdot \pi_{V_{3}}^{V_{2}}\left(v_{2}\right)+\gamma_{V_{3}}\left(\neg v_{3} \mid v_{1} \wedge \neg v_{2}\right) \cdot \pi_{V_{3}}^{V_{2}}\left(\neg v_{2}\right)\right] \\
= & \gamma_{V_{3}}\left(v_{3} \mid v_{1} \wedge v_{2}\right) \cdot \pi_{V_{3}}^{V_{2}}\left(v_{2}\right)+\gamma_{V_{3}}\left(v_{3} \mid v_{1} \wedge \neg v_{2}\right) \cdot \pi_{V_{3}}^{V_{2}}\left(\neg v_{2}\right)= \\
= & 0.75 \cdot 0.5+0.25 \cdot 0.5=0.5 \\
\lambda_{V_{3}}^{V_{1}}\left(\neg v_{1}\right)= & \lambda_{V_{3}}\left(v_{3}\right) \cdot\left[\gamma_{V_{3}}\left(v_{3} \mid \neg v_{1} \wedge v_{2}\right) \cdot \pi_{V_{3}}^{V_{2}}\left(v_{2}\right)+\gamma_{V_{3}}\left(v_{3} \mid \neg v_{1} \wedge \neg v_{2}\right) \cdot \pi_{V_{3}}^{V_{2}}\left(\neg v_{2}\right)\right] \\
& +\lambda_{V_{3}}\left(\neg v_{3}\right) \cdot\left[\gamma_{V_{3}}\left(\neg v_{3} \mid \neg v_{1} \wedge v_{2}\right) \cdot \pi_{V_{3}}^{V_{2}}\left(v_{2}\right)+\gamma_{V_{3}}\left(\neg v_{3} \mid \neg v_{1} \wedge \neg v_{2}\right) \cdot \pi_{V_{3}}^{V_{2}}\left(\neg v_{2}\right)\right] \\
= & \gamma_{V_{3}}\left(v_{3} \mid \neg v_{1} \wedge v_{2}\right) \cdot \pi_{V_{3}}^{V_{2}}\left(v_{2}\right)+\gamma_{V_{3}}\left(v_{3} \mid \neg v_{1} \wedge \neg v_{2}\right) \cdot \pi_{V_{3}}^{V_{2}}\left(\neg v_{2}\right) \\
= & 0.4 \cdot 0.5+0.2 \cdot 0.5=0.3
\end{aligned}
$$

From the diagnostic parameter $\lambda_{V_{3}}^{V_{1}}$ it receives from its successor $V_{3}$, vertex $V_{1}$ now finds the values of its compound diagnostic parameter to be

$$
\begin{array}{ll}
\lambda_{V_{1}}\left(v_{1}\right) & =0.5 \\
\lambda_{V_{1}}\left(\neg v_{1}\right) & =0.3
\end{array}
$$

Substitution of the values of the compound parameters into the data fusion lemma yields

$$
\begin{array}{ll}
\operatorname{Pr}^{v_{3}}\left(v_{1}\right) & =\alpha \cdot 0.25 \cdot 0.5=\alpha \cdot 0.125 \\
\operatorname{Pr}^{v_{3}}\left(\neg v_{1}\right) & =\alpha \cdot 0.75 \cdot 0.3=\alpha \cdot 0.225
\end{array}
$$

After elimination of the normalisation constant $\alpha$ (note that we postponed all normalisation until this last step!), vertex $V_{1}$ finds

$$
\begin{array}{ll}
\operatorname{Pr}^{v_{3}}\left(v_{1}\right) & =0.357 \\
\operatorname{Pr}^{v_{3}}\left(\neg v_{1}\right) & =0.643
\end{array}
$$

To conclude our discussion of Pearl's algorithm so far, we take a (superficial) look at the algorithm's computational complexity. An analysis of the computation rules involved in the algorithm serves to show that it has a worst-case computational time complexity that is exponential in the number of vertices in a network's digraph. We consider, in a Bayesian network with $n$ vertices, a vertex $V_{i}$ with $O(n)$ predecessors and $O(n)$ successors. For this vertex, computing the compound causal parameter requires at most $O\left(2^{n}\right)$ computations; computing its compound diagnostic parameter requires at most $O(n)$ computations. Vertex $V_{i}$ can therefore compute the probabilities of its values in at most $O\left(2^{n}\right)$ time. The computation of a single diagnostic parameter from $V_{i}$ requires at most $O\left(2^{n}\right)$ computations; the computation of a single causal parameter requires only constant time. Vertex $V_{i}$ can therefore compute all parameters to send to its various neighbours in at most $O\left(n \cdot 2^{n}\right)$ time. Propagating a single piece of evidence throughout the network may thus require $O\left(n^{2} \cdot 2^{n}\right)$ computations. The exponential factor in the computational time complexity of Pearl's algorithm arises solely from the number of predecessors a vertex can have in a network's digraph: if the predecessor sets are bounded in size by a constant, then the algorithm has a linear runtime complexity.

### 4.2.3 Multiply Connected Digraphs

In the previous section, we have detailed Pearl's algorithm for probabilistic inference with a Bayesian network comprising a singly connected digraph for its qualitative part. We will now extend the algorithm to apply to Bayesian networks of which the qualitative part is a multiply connected digraph.

We recall from our discussion of Pearl's algorithm so far, that the computation rules for probabilistic inference with a Bayesian network comprising a singly connected digraph derive from exploiting independences that are read from the network's graph by local inspection of a vertex' incoming and outgoing arcs. These computation rules therefore make use explicitly of the property that in a singly connected digraph there is at most one chain between any two vertices. We observe that this property does not hold in a multiply connected digraph as in such a graph there may be multiple chains between vertices. Local inspection of a vertex' incident arcs now no longer suffices for reading independences from a Bayesian network's digraph. In fact, application of Pearl's algorithm discussed so far to a Bayesian network comprising a multiply connected digraph inevitably leads to various problems due to the presence of loops in the network's digraph [Suermondt \& Cooper, 1990]: vertices may indefinitely send newly updated messages, originating from the same evidence, to their neighbours, causing the network never to reach a new equilibrium, and even if the network does reach an equilibrium, it is not guaranteed to reflect the correct updated joint probability distribution. Pearl's algorithm, therefore, cannot be extended based on similar ideas as before to apply to networks with a multiply connected digraph: the algorithm has to be supplemented with an additional method for coping with the loops in such a digraph.

The best-known method for coping with loops in probabilistic inference is the method of loop cutset conditioning [Pearl, 1988]. The idea underlying the method of loop cutset conditioning is that of reasoning by assumption. For a multiply connected digraph of a Bayesian network, vertices are selected that, upon instantiation, with each other effectively 'cut' all loops of the digraph and cause it to 'behave' as if it were singly connected; the selected vertices are said to constitute a loop cutset for the network's digraph. During probabilistic inference with the network, each configu-
ration of the selected loop cutset is looked upon as a (compound) assumption on which reasoning is performed: for each vertex, the probabilities of its values are computed by conditioning successively on all possible configurations of the loop cutset and subsequently weighting the results obtained with the probabilities of these configurations. Since upon instantiation of the loop cutset the digraph behaves as if it were singly connected, the probabilities of a vertex' values can simply be computed from the network using Pearl's basic algorithm for Bayesian networks with a singly connected digraph.

## Loop Cutsets

We begin our discussion of the method of loop cutset conditioning by detailing the concept of a loop cutset for a Bayesian network's multiply connected digraph. A loop cutset has to provide, upon instantiation, for the multiply connected digraph to behave as if it were singly connected, regardless of the evidence that has been entered into the network at hand. Note that a loop cutset then serves to prevent, at any time, probabilistic information originating at a single vertex in the digraph to reach another vertex along multiple chains. For this purpose, a loop cutset has to cut, or block, every cyclic chain, or loop, in the network's digraph. Note that it does not suffice to block a cyclic chain by the empty set as this chain may become unblocked as further evidence is entered into the network. From these observations, we find that a loop cutset has to block all cyclic chains in the digraph in such a way that each such chain contains at least one vertex from the loop cutset that has an outgoing arc on the chain.

Definition 4.2.16 Let $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ be an acyclic multiply connected digraph. A set of vertices $\boldsymbol{L}_{G} \subseteq \boldsymbol{V}_{G}$ is called a loop cutset for $G$ if each loop $s$ in $G$ contains three consecutive vertices $X_{1}, X_{2}, X_{3}$, for which one of the following conditions holds:

- arcs $\left(X_{2}, X_{1}\right)$ and $\left(X_{2}, X_{3}\right)$ are on the loop $s$, and $X_{2} \in \boldsymbol{L}_{G}$;
- arcs $\left(X_{1}, X_{2}\right)$ and $\left(X_{2}, X_{3}\right)$ are on the loop $s$, and $X_{2} \in \boldsymbol{L}_{G}$.

Note that a Bayesian network's multiply connected digraph $G$ may allow several different loop cutsets; in fact, any superset of a loop cutset for $G$ is also a loop cutset for $G$. Also note that any loop cutset for $G$ is a non-empty set of vertices.

Example 4.2.17 We consider the multiply connected digraph $G$ shown in Figure 4.9. The sets of vertices $\left\{V_{2}, V_{6}, V_{9}\right\},\left\{V_{3}, V_{4}, V_{9}\right\}$, and $\left\{V_{2}, V_{8}, V_{10}, V_{11}\right\}$ are example loop cutsets for $G$.

## Loop Cutset Conditioning

After having detailed the concept of a loop cutset, we now turn to the method of loop cutset conditioning itself. To this end, we consider a Bayesian network $\mathcal{B}=(G, \Gamma)$ where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is an acyclic multiply connected digraph; let Pr be the joint probability distribution defined by $\mathcal{B}$. Furthermore, suppose that we have the loop cutset $\boldsymbol{L}_{G}$ for the digraph $G$ of the network, as defined above. We now address computing the probabilities of the values of some vertex $V_{i}$ in $G$, that is, we consider the computation of the probabilities $\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{G}}\right)$. As the digraph $G$ is multiply connected, these probabilities cannot be computed directly from the network using Pearl's algorithm. Since the loop cutset has been chosen so as to block, upon instantiation, all cyclic chains in


Figure 4.9: An Example Multiply Connected Digraph.
the digraph, we observe that any probability that is conditioned on a configuration of (at least)the loop cutset can be computed from the network directly. Now, by conditioning the probabilities $\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{G}}\right)$ of interest on the various configurations of the loop cutset $\boldsymbol{L}_{G}$ for $G$, we find

$$
\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{G}}\right)=\sum_{c_{\boldsymbol{L}_{G}}} \operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{G}} \wedge c_{\boldsymbol{L}_{G}}\right) \cdot \operatorname{Pr}\left(c_{\boldsymbol{L}_{G}} \mid \widetilde{c}_{\boldsymbol{V}_{G}}\right)
$$

The probabilities $\operatorname{Pr}\left(V_{i} \mid \widetilde{c}_{V_{G}} \wedge c_{\boldsymbol{L}_{G}}\right)$ are now conditioned on a configuration of (at least) the loop cutset and therefore can be computed directly from the Bayesian network using Pearl's basic algorithm. The probabilities $\operatorname{Pr}\left(c_{\boldsymbol{L}_{G}} \mid \widetilde{c}_{V_{G}}\right)$, however, cannot be computed from the Bayesian network directly. For computing these probabilities, the following recursive computation rules are used:

- the prior probabilities $\operatorname{Pr}\left(c_{\boldsymbol{L}_{G}}\right)$ are computed from the joint probability distribution Pr using the property stated in Proposition 4.1.3 [Suermondt \& Cooper, 1991].
- for the first piece of evidence $e_{1}$ entered into the Bayesian network, the (updated) probabilities of the configurations of the loop cutset are computed from

$$
\operatorname{Pr}\left(c_{\boldsymbol{L}_{G}} \mid e_{1}\right)=\alpha \cdot \operatorname{Pr}\left(e_{1} \mid c_{\boldsymbol{L}_{G}}\right) \cdot \operatorname{Pr}\left(c_{\boldsymbol{L}_{G}}\right)
$$

where $\alpha$ is a normalisation constant. The probabilities $\operatorname{Pr}\left(e_{1} \mid c_{\boldsymbol{L}_{G}}\right)$ are computed directly from the network using Pearl's basic algorithm for singly connected digraphs; the probabilities $\operatorname{Pr}\left(c_{\boldsymbol{L}_{G}}\right)$ have been computed before and, hence, are already available.

- for the $j$-th piece of evidence $e_{j}$ entered into the network, the (updated) probabilities of the configurations of the loop cutset are computed from

$$
\begin{aligned}
\operatorname{Pr}\left(c_{\boldsymbol{L}_{G}} \mid e_{1} \wedge \cdots \wedge e_{j}\right)= & \alpha \cdot \operatorname{Pr}\left(e_{j} \mid c_{\boldsymbol{L}_{G}} \wedge e_{1} \wedge \cdots \wedge e_{j-1}\right) \\
& \cdot \operatorname{Pr}\left(c_{\boldsymbol{L}_{G}} \mid e_{1} \wedge \cdots \wedge e_{j-1}\right)
\end{aligned}
$$

where $\alpha$ once more is a normalisation constant. The probabilities $\operatorname{Pr}\left(e_{j} \mid c_{\boldsymbol{L}_{G}} \wedge\right.$ $e_{1} \wedge \cdots \wedge e_{j-1}$ ) are computed directly from the Bayesian network using Pearl's basic algorithm; the probability $\operatorname{Pr}\left(c_{\boldsymbol{L}_{G}} \mid e_{1} \wedge \cdots \wedge e_{j-1}\right)$ has been computed before in the previous step of the recursion and need not be re-computed.

Note that the probabilities of the configurations of a digraph's loop cutset involve information from the probability distribution concerning various vertices combined. These probabilities therefore cannot be computed from the network by local message passing between neighbouring vertices alone. For this purpose, the computational architecture is enhanced with a global supervisor that is capable of performing the recursive computation rules outlined above.

We briefly illustrate the method of loop cutset conditioning by means of an example.
Example 4.2.18 We consider the Bayesian network $\mathcal{B}=(G, \Gamma)$ shown in Figure 4.10. Let $\operatorname{Pr}$ be the joint probability distribution defined by $\mathcal{B}$. We address the computation of the probabilities $\operatorname{Pr}\left(v_{4}\right)$ and $\operatorname{Pr}\left(\neg v_{4}\right)$ using Pearl's algorithm.


Figure 4.10: An Example Bayesian Network.
We begin by observing that the digraph $G$ of the Bayesian network is multiply connected. To compute the probabilities of interest from the network, we enhance Pearl's algorithm with the method of loop cutset conditioning. Suppose that for this purpose we choose for the digraph $G$ the loop cutset $\boldsymbol{L}_{G}=\left\{V_{2}\right\}$. The global supervisor of the architecture now starts the computation by computing the probabilities of the configurations of the loop cutset. To this end, it employs the property stated in Proposition 4.1.3 to find that

$$
\begin{aligned}
\operatorname{Pr}\left(v_{2}\right) & =\operatorname{Pr}\left(v_{1} \wedge v_{2}\right)+\operatorname{Pr}\left(\neg v_{1} \wedge v_{2}\right)= \\
& =\gamma_{V_{2}}\left(v_{2} \mid v_{1}\right) \cdot \gamma_{V_{1}}\left(v_{1}\right)+\gamma_{V_{2}}\left(v_{2} \mid \neg v_{1}\right) \cdot \operatorname{Pr}\left(\neg v_{1}\right)= \\
& =0.1 \cdot 0.6+0.35 \cdot 0.4=0.2 \\
\operatorname{Pr}\left(\neg v_{2}\right) & =\operatorname{Pr}\left(v_{1} \wedge \neg v_{2}\right)+\operatorname{Pr}\left(\neg v_{1} \wedge \neg v_{2}\right)= \\
& =\gamma_{V_{2}}\left(\neg v_{2} \mid v_{1}\right) \cdot \gamma_{V_{1}}\left(v_{1}\right)+\gamma_{V_{2}}\left(\neg v_{2} \mid \neg v_{1}\right) \cdot \operatorname{Pr}\left(\neg v_{1}\right)= \\
& =0.9 \cdot 0.6+0.65 \cdot 0.4=0.8
\end{aligned}
$$

The probabilities of interest are now computed by conditioning on the different configurations of the chosen loop cutset:

$$
\begin{array}{ll}
\operatorname{Pr}\left(v_{4}\right) & =\operatorname{Pr}\left(v_{4} \mid v_{2}\right) \cdot \operatorname{Pr}\left(v_{2}\right)+\operatorname{Pr}\left(v_{4} \mid \neg v_{2}\right) \cdot \operatorname{Pr}\left(\neg v_{2}\right) \\
\operatorname{Pr}\left(\neg v_{4}\right) & =\operatorname{Pr}\left(\neg v_{4} \mid v_{2}\right) \cdot \operatorname{Pr}\left(v_{2}\right)+\operatorname{Pr}\left(\neg v_{4} \mid \neg v_{2}\right) \cdot \operatorname{Pr}\left(\neg v_{2}\right)
\end{array}
$$

The global supervisor provisionally enters the value true for vertex $V_{2}$ into the network to enable vertex $V_{4}$ to compute the probabilities $\operatorname{Pr}\left(v_{4} \mid v_{2}\right)$ and $\operatorname{Pr}\left(\neg v_{4} \mid v_{2}\right)$. Vertex $V_{4}$ computes these probabilities to be

$$
\begin{array}{ll}
\operatorname{Pr}\left(v_{4} \mid v_{2}\right) & =0.6 \\
\operatorname{Pr}\left(\neg v_{4} \mid v_{2}\right) & =0.4
\end{array}
$$

The global supervisor now retracts the value true for vertex $V_{2}$ and subsequently enters the value false for the vertex. It now requests vertex $V_{4}$ to compute the probabilities $\operatorname{Pr}\left(v_{4} \mid \neg v_{2}\right)$ and $\operatorname{Pr}\left(\neg v_{4} \mid \neg v_{2}\right)$. Vertex $V_{4}$ returns

$$
\begin{array}{ll}
\operatorname{Pr}\left(v_{4} \mid \neg v_{2}\right) & =0.1 \\
\operatorname{Pr}\left(\neg v_{4} \mid \neg v_{2}\right) & =0.9
\end{array}
$$

The global supervisor substitutes the probabilities returned by vertex $V_{4}$ and the probabilities it has itself computed for the loop cutset, into the conditioning formula to find

$$
\begin{array}{ll}
\operatorname{Pr}\left(v_{4}\right) & =0.6 \cdot 0.2+0.1 \cdot 0.8=0.2 \\
\operatorname{Pr}\left(\neg v_{4}\right) & =0.4 \cdot 0.2+0.9 \cdot 0.8=0.8
\end{array}
$$

Now suppose that the evidence $V_{3}=$ true is observed and entered into the Bayesian network $\mathcal{B}$. We address the computation of the updated probabilities of the values of vertex $V_{4}$. The global supervisor begins with updating the probabilities of the configurations of the loop cutset:

$$
\begin{array}{ll}
\operatorname{Pr}\left(v_{2} \mid v_{3}\right) & =\alpha \cdot \operatorname{Pr}\left(v_{3} \mid v_{2}\right) \cdot \operatorname{Pr}\left(v_{2}\right) \\
\operatorname{Pr}\left(\neg v_{2} \mid v_{3}\right) & =\alpha \cdot \operatorname{Pr}\left(v_{3} \mid \neg v_{2}\right) \cdot \operatorname{Pr}\left(\neg v_{2}\right)
\end{array}
$$

where $\alpha$ is a normalisation constant. The supervisor requests from vertex $V_{3}$ the probabilities $\operatorname{Pr}\left(v_{3} \mid v_{2}\right)$ and $\operatorname{Pr}\left(v_{3} \mid \neg v_{2}\right)$; to enable vertex $V_{3}$ to compute these probabilities, it provisionally enters and retracts the values true and false for vertex $V_{2}$, respectively. Vertex $V_{3}$ computes the requested probabilities to be

$$
\begin{array}{ll}
\operatorname{Pr}\left(v_{3} \mid v_{2}\right) & =0.9 \\
\operatorname{Pr}\left(v_{3} \mid \neg v_{2}\right) & =0.4
\end{array}
$$

After substitution and subsequent elimination of the normalisation constant $\alpha$, the global supervisor finds

$$
\begin{array}{ll}
\operatorname{Pr}\left(v_{2} \mid v_{3}\right) & =0.36 \\
\operatorname{Pr}\left(\neg v_{2} \mid v_{3}\right) & =0.64
\end{array}
$$

The probabilities of interest are now computed by conditioning as before:

$$
\begin{aligned}
\operatorname{Pr}^{v_{3}}\left(v_{4}\right) & =\operatorname{Pr}\left(v_{4} \mid v_{2} \wedge v_{3}\right) \cdot \operatorname{Pr}\left(v_{2} \mid v_{3}\right)+\operatorname{Pr}\left(v_{4} \mid \neg v_{2} \wedge v_{3}\right) \cdot \operatorname{Pr}\left(\neg v_{2} \mid v_{3}\right)= \\
& =0.6 \cdot 0.36+0.1 \cdot 0.64=0.28 \\
\operatorname{Pr}^{v_{3}}\left(\neg v_{4}\right) & =\operatorname{Pr}\left(\neg v_{4} \mid v_{2} \wedge v_{3}\right) \cdot \operatorname{Pr}\left(v_{2} \mid v_{3}\right)+\operatorname{Pr}\left(\neg v_{4} \mid \neg v_{2} \wedge v_{3}\right) \cdot \operatorname{Pr}\left(\neg v_{2} \mid v_{3}\right)= \\
& =0.4 \cdot 0.36+0.9 \cdot 0.64=0.72
\end{aligned}
$$

To conclude our discussion of the enhancement of Pearl's algorithm for probabilistic inference with the method of loop cutset conditioning, we take a (superficial) look at the algorithm's computational complexity. We recall that analysis of Pearl's algorithm for Bayesian networks comprising a singly connected digraph has revealed that a vertex can compute the probabilities of its values in at most $O\left(2^{n}\right)$ time, where $n$ is the number of vertices in the network's digraph. Propagating a piece of evidence throughout the network may require as many as $O\left(n^{2} \cdot 2^{n}\right)$ computations. An analysis of the method of loop cutset conditioning now serves to show that the method adds a factor to the basic algorithm's computational complexity that is exponential in the size of the loop cutset employed. Initialising the probabilities for the loop cutset's configurations requires $O\left(2^{l}\right)$ computations, where $l$ is the number of vertices in the loop cutset. Computing the probabilities of a vertex' values now takes $O\left(2^{n+l}\right)$ time, since these probabilities have to be computed conditioned on every possible configuration of the loop cutset. Propagating a piece of evidence requires at most $O\left(n^{2} \cdot 2^{n+l}\right)$ computations.

## Finding a Loop Cutset

In the foregoing, we have argued that the computational complexity of Pearl's algorithm enhanced with loop cutset conditioning, for probabilistic inference with Bayesian networks comprising a multiply connected digraph, is exponential in the size of the loop cutset employed. From a computational point of view, therefore, the best loop cutset to use in practical applications is a loop cutset having the smallest number of vertices. A loop cutset with the smallest number of vertices is called an optimal loop cutset; we define the concept of an optimal loop cutset more formally.

Definition 4.2.19 Let $G$ be an acyclic multiply connected digraph. A loop cutset $\boldsymbol{L}_{G}$ for $G$ is called minimal for $G$ if no proper subset of $\boldsymbol{L}_{G}$ is a loop cutset for $G$; the loop cutset $\boldsymbol{L}_{G}$ is called optimal for $G$ if for any other loop cutset $\boldsymbol{L}_{G}^{\prime}$ for $G$ we have that $\left|\boldsymbol{L}_{G}\right| \leq\left|\boldsymbol{L}_{G}^{\prime}\right|$.

Note that an optimal loop cutset for a multiply connected digraph also is a minimal loop cutset for this digraph; the reverse in general is not true. Also note that an optimal loop cutset for a digraph need not be unique.

Example 4.2.20 We consider once more the multiply connected digraph $G$ shown in Figure 4.9. The sets of vertices $\left\{V_{2}, V_{6}, V_{9}\right\}$ and $\left\{V_{3}, V_{4}, V_{9}\right\}$ are optimal loop cutsets for $G$. The set of vertices $\left\{V_{2}, V_{8}, V_{10}, V_{11}\right\}$ is a minimal loop cutset for $G$ that is not optimal. The set of vertices $\left\{V_{2}, V_{4}, V_{8}, V_{9}\right\}$ is a loop cutset that is neither optimal nor minimal for $G$.

Unfortunately, the problem of finding for an acyclic multiply connected digraph an optimal loop cutset is known to be NP-hard [Suermondt \& Cooper, 1990]; hence, it is not likely that a generally applicable polynomial-time algorithm for this problem will be found. For the purpose of finding a good loop cutset for a given multiply connected digraph, therefore, generally a heuristic algorithm is used. The best known among these algorithms is a heuristic algorithm designed by H.J. Suermondt and G.F. Cooper [Suermondt \& Cooper, 1990].

The loop-cutset algorithm of Suermondt and Cooper takes an acyclic multiply connected digraph $G$ for its input and yields a loop cutset $\boldsymbol{L}_{G}$ for $G$ as its output. The
algorithm is composed of two basic phases that are alternated recursively. In the first phase, the algorithm deletes from the digraph $G$ as many vertices as possible that are not included in any loop; such vertices need not be considered as candidates for the loop cutset as upon instantiation they will never block any loops in $G$. For this purpose, all vertices that have at most one incident arc are removed recursively from $G$, along with their incident arcs. The second phase of the algorithm is entered with a reduced digraph $G^{\prime}$ in which each vertex has at least two incident arcs. From among this digraph's vertices, a single vertex is selected for inclusion in the loop cutset in the making. The selection of this vertex is based on a heuristic criterion that aims at selecting a vertex that upon instantiation blocks as many loops from the reduced digraph $G^{\prime}$ as possible. This heuristic criterion is based on the idea that a vertex with a larger number of incident arcs in $G^{\prime}$ is likely to be included in more loops than a vertex with a smaller number of incident arcs; however, a vertex with two incoming arcs on a loop does not serve to block this loop. The heuristic criterion therefore selects a vertex $V_{i}$ with highest degree having at most one predecessor. The thus selected vertex is included in the loop cutset in the making. The vertex $V_{i}$ is subsequently removed from the digraph, along with all its incident arcs. Note that in the thus reduced digraph all loops that are blocked by vertex $V_{i}$ are effectively cut, leaving only loops that are yet unblocked. The algorithm then returns to the first phase. The two phases of the algorithm are thus alternated recursively until no more vertices remain to be considered in the reduced graph. Figure 4.11 summarises the heuristic loop-cutset algorithm of Suermondt and Cooper in pseudo-code. We illustrate the heuristic loop-cutset

```
procedure loop-cutset \(\left(G, \boldsymbol{L}_{G}\right)\)
    \(\boldsymbol{L}_{G}:=\varnothing ;\)
    while \(V_{G} \neq \varnothing\) do
        if there is a vertex \(V \in \boldsymbol{V}_{G}\) with \(\operatorname{degree}(V) \leq 1\)
        then select vertex \(V\)
        else \(\boldsymbol{K}:=\left\{V_{i} \mid V_{i} \in \boldsymbol{V}_{G}\right.\) and \(V_{i}\) has indegree \(\left.\left(V_{i}\right) \leq 1\right\}\);
            select from \(\boldsymbol{K}\) a vertex \(V\) with \(\operatorname{degree}(V) \geq \operatorname{degree}\left(V_{i}\right)\) for all \(V_{i} \in \boldsymbol{K}\);
            \(\boldsymbol{L}_{G}:=\boldsymbol{L}_{G} \cup\{V\}\)
        fi;
        \(\boldsymbol{V}_{G}:=\boldsymbol{V}_{G} \backslash\{V\} ;\)
        \(\boldsymbol{A}_{G}:=\boldsymbol{A}_{G} \cap\left(\boldsymbol{V}_{G} \times \boldsymbol{V}_{G}\right)\)
    od
end
```

Figure 4.11: The Loop-cutset Algorithm of Suermondt and Cooper.
algorithm outlined above by means of an example.
Example 4.2.21 We consider once again the multiply connected digraph $G$ from Figure 4.9 and address the construction of a loop cutset for $G$ using Suermondt and Cooper's heuristic algorithm.

The algorithm sets out by initialising the loop cutset in the making with the empty set:

$$
\boldsymbol{L}_{G}=\varnothing
$$


(a)

(b)

(c)

Figure 4.12: Application of the Loop-cutset Algorithm.

The algorithm now notes that the digraph $G$ comprises a single vertex with at most one incident arc - this is the vertex $V_{1}$. The algorithm selects this vertex and removes it, along with its incident arc, from $G$. The thus reduced digraph $G_{1}$ is shown in Figure 4.12 (a). From the digraph $G_{1}$, the algorithm selects all vertices with at most one incoming arc. These vertices constitute the set of candidates for inclusion in the loop cutset in the making:

$$
\boldsymbol{K}_{1}=\left\{V_{2}, V_{3}, V_{4}, V_{5}, V_{6}, V_{8}, V_{10}, V_{11}\right\}
$$

From among these candidates, the algorithm selects a vertex with highest degree. We suppose that the algorithm selects vertex $V_{2}$ to include it in the loop cutset in the making:

$$
\boldsymbol{L}_{G}=\left\{V_{2}\right\}
$$

The algorithm proceeds by removing from $G_{1}$ the selected vertex along with its incident arcs. After subsequently removing all vertices with at most one incident arc, the reduced
digraph $G_{2}$ shown in Figure 4.12 (b) results. From the digraph $G_{2}$, the algorithm once more selects the set of candidates for inclusion in the loop cutset:

$$
\boldsymbol{K}_{\mathbf{2}}=\left\{V_{4}, V_{6}, V_{7}, V_{8}, V_{10}, V_{11}\right\}
$$

From among these candidates, the algorithm selects a vertex with highest degree. We suppose that it selects vertex $V_{6}$. Vertex $V_{6}$ is now added to the loop cutset in the making:

$$
\boldsymbol{L}_{G}=\left\{V_{2}, V_{6}\right\}
$$

After removal of vertex $V_{6}$ and its incident arcs from $G_{2}$, and recursive removal of all vertices with at most one incident arc, the reduced digraph $G_{3}$ shown in Figure 4.12 (c) is yielded. From the digraph $G_{3}$, the algorithm selects the set of candidates to be

$$
\boldsymbol{K}_{\mathbf{3}}=\left\{V_{9}, V_{10}, V_{11}\right\}
$$

From among these candidates, the algorithm selects vertex $V_{9}$ for inclusion in the loop cutset as it has the highest degree:

$$
\boldsymbol{L}_{G}=\left\{V_{2}, V_{6}, V_{9}\right\}
$$

After removing vertex $V_{9}$ with its incident arcs from $G_{3}$, subsequent removal of all vertices with at most one incident arc results in the empty graph. The algorithm halts and yields the loop cutset $L_{G}=\left\{V_{2}, V_{6}, V_{9}\right\}$ for the digraph $G$. Note that the loop cutset yielded by the algorithm is an optimal loop cutset for $G$.

The heuristic loop-cutset algorithm of Suermondt and Cooper is correct in the sense that, for any acyclic multiply connected digraph, the algorithm will halt and yield a set of vertices that is a loop cutset of the digraph [Suermondt \& Cooper, 1990]. The heuristic algorithm, however, is not guaranteed to always yield an optimal loop cutset. Experimental results have shown that for randomly generated acyclic multiply connected digraphs, the algorithm finds an optimal loop cutset in approximately $70 \%$ of the studied digraphs [Suermondt \& Cooper, 1990].

### 4.2.4 Other Algorithms for Probabilistic Inference

In the previous sections, we have discussed Pearl's basic algorithm and its enhancement for general probabilistic inference with a Bayesian network. Pearl's algorithm, however, is not the only algorithm designed for this purpose: several different algorithms have been proposed in the course of the preceding decade. In general, an algorithm for probabilistic inference with a Bayesian network is built from two basic procedures:

- a procedure for (efficiently) computing probabilities of interest from a Bayesian network, and
- a procedure for processing evidence, that is, for entering evidence into the network and subsequently (efficiently) computing the updated joint probability distribution given the evidence.

Note that in Pearl's algorithm these basic procedures are combined into a single set of computation rules where they cannot easily be distinguished. In most other algorithms, however, these basic procedures are more readily discernible. The algorithms proposed further have two important properties in common: the qualitative part of a Bayesian network is exploited more or less directly as a computational architecture, and after a piece of evidence has been processed in the network, again a Bayesian network results. Note that the latter property allows for recursive processing of evidence.

Although all algorithms proposed for probabilistic inference build on the same notion of a Bayesian network, they differ considerably with respect to their underlying concepts. To support this observation, we briefly review another algorithm for probabilistic inference with a Bayesian network, designed by S.L. Lauritzen and D.J. Spiegelhalter [Lauritzen \& Spiegelhalter, 1988]. Lauritzen and Spiegelhalter have observed that, after a piece of evidence has become available, updating the joint probability distribution generally entails going against the 'direction' of the initially assessed conditional probabilities. From this observation they have concluded that the directed qualitative part of a Bayesian network is not suitable as an architecture for probabilistic inference directly. Their algorithm therefore builds on an undirected representation of a joint probability distribution. Prior to probabilistic inference, the original Bayesian network is transformed into a so-called decomposable probabilistic network. A decomposable probabilistic network again consists of a qualitative part and a quantitative part. The qualitative part is a decomposable, or chordal, graph; the quantitative part is a set of marginal distributions on the vertex sets of the cliques of this graph. The computational architecture for the algorithm derives from the qualitative part of this decomposable probabilistic network in the following sense: the cliques of the decomposable graph are the autonomous objects in the architecture, and the clique intersections give rise to its communication channels. From a decomposable probabilistic network, a probability of interest is computed by further marginalisation of an appropriate marginal distribution. Processing evidence is performed per clique and by message-passing between neighbouring cliques in the architecture.

The various algorithms proposed for probabilistic inference with a Bayesian network also differ with respect to their computational complexity. We note that probabilistic inference with Bayesian networks without any topological restrictions is known to be NP-hard [Cooper, 1990]. All algorithms proposed therefore have an exponential worst-case computational complexity. However, the algorithms can be shown to behave polynomially under certain restrictions. These restrictions concern the topology of a Bayesian network's digraph and differ among the various algorithms. In general, the sparser the digraph of a Bayesian network, the better most algorithms perform. Experience with constructing Bayesian networks for real-life applications so far has indicated that in fact a Bayesian network's digraph tends to be sparse.

## Exercises

## Exercise 4.1

Consider the following acyclic digraph $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ :


Now suppose that the graph $G$ is a (minimal) directed I-map for the independence relation of the joint probability distribution $\operatorname{Pr}$ on $\boldsymbol{V}_{G}$. State which (conditional) probabilities from the distribution $\operatorname{Pr}$ have to be associated with $G$ to arrive at a Bayesian network comprising $G$ for its qualitative part.

## * Exercise 4.2

Consider the Bayesian network $\mathcal{B}=(G, \Gamma)$ where $G$ is the following digraph:

and $\Gamma=\left\{\gamma_{V_{i}} \mid V_{i} \in \boldsymbol{V}_{G}\right\}$ is defined as:

$$
\begin{array}{llll}
\gamma_{V_{1}}\left(v_{1}\right) & =0.25 & \gamma_{V_{2}}\left(v_{2}\right) & =0.5 \\
\gamma_{V_{3}}\left(v_{3} \mid v_{1} \wedge v_{2}\right) & =0.75 & & \gamma_{V_{4}}\left(v_{4} \mid v_{3}\right) \\
\gamma_{V_{3}}\left(v_{3} \mid \neg v_{1} \wedge v_{2}\right) & =0.4 & & =0.8 \\
\gamma_{V_{3}}\left(v_{3} \mid v_{1} \wedge \neg v_{2}\right) & =0.25 & & \\
\gamma_{V_{3}}\left(v_{3} \mid \neg v_{1} \wedge \neg v_{3}\right) & =0 \\
& =0.2 & &
\end{array}
$$

Let $\operatorname{Pr}$ be the joint probability distribution defined by $\mathcal{B}$. Compute the following probabilities by exploiting the property stated in Proposition 4.1.3:
a. $\operatorname{Pr}\left(v_{1} \wedge v_{2} \wedge v_{3}\right) ;$
b. $\operatorname{Pr}\left(v_{2} \wedge v_{3}\right)$;
c. $\operatorname{Pr}\left(v_{1} \mid v_{2} \wedge v_{3}\right) ;$
d. $\operatorname{Pr}\left(v_{1} \vee v_{2} \vee \neg v_{3} \vee v_{4}\right)$.

## * Exercise 4.3

Let $\mathcal{B}=(G, \Gamma)$ be a Bayesian network where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is a directed tree. Suppose that Pearl's algorithm is employed for probabilistic inference with $\mathcal{B}$. For each vertex $V_{i} \in \boldsymbol{V}_{G}$, let $\boldsymbol{V}_{\boldsymbol{i}}^{-}=\boldsymbol{\sigma}_{G}^{*}\left(V_{i}\right)$; let $\lambda_{V_{i}}\left(V_{i}\right)$ be the compound diagnostic parameter for $V_{i}$. Show that if $\widetilde{c}_{\boldsymbol{V}_{\boldsymbol{i}}^{-}}=$True, then $\lambda_{V_{i}}\left(V_{i}\right)=1$, for all $V_{i} \in \boldsymbol{V}_{G}$.

## * Exercise 4.4

Consider the Bayesian network $\mathcal{B}=(G, \Gamma)$ where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is the following digraph:

and $\Gamma=\left\{\gamma_{V_{i}} \mid V_{i} \in \boldsymbol{V}_{G}\right\}$ is defined as

$$
\begin{array}{ll}
\gamma_{V_{1}}\left(v_{1}\right)=0.5 & \gamma_{V_{2}}\left(v_{2}\right)=0.4 \\
\gamma_{V_{3}}\left(v_{3} \mid v_{1} \wedge v_{2}\right)=0.9 & \gamma_{V_{4}}\left(v_{4} \mid v_{3}\right)=0.5 \\
\gamma_{V_{3}}\left(v_{3} \mid \neg v_{1} \wedge v_{2}\right)=0.6 & \gamma_{V_{4}}\left(v_{4} \mid \neg v_{3}\right)=0.65 \\
\gamma_{V_{3}}\left(v_{3} \mid v_{1} \wedge \neg v_{2}\right)=0.2 & \\
\gamma_{V_{3}}\left(v_{3} \mid \neg v_{1} \wedge \neg v_{2}\right)=0.8 & \\
& \\
\gamma_{V_{5}}\left(v_{5} \mid v_{4}\right)=0.25 & \gamma_{V_{6}}\left(v_{6} \mid v_{4}\right)=0.6 \\
\gamma_{V_{5}}\left(v_{5} \mid \neg v_{4}\right)=0.3 & \gamma_{V_{6}}\left(v_{6} \mid \neg v_{4}\right)=0.1
\end{array}
$$

Let $\operatorname{Pr}$ be the joint probability distribution defined by the Bayesian network $\mathcal{B}$. Suppose that Pearl's algorithm is employed for probabilistic inference with $\mathcal{B}$.
a. Compute the probabilities $\operatorname{Pr}\left(v_{4}\right)$ and $\operatorname{Pr}\left(\neg v_{4}\right)$ for vertex $V_{4}$ from the Bayesian network $\mathcal{B}$.
b. Suppose that the evidence $V_{6}=$ true is observed and entered into the network. Compute the updated probabilities $\operatorname{Pr}^{v_{6}}\left(v_{4}\right)$ and $\operatorname{Pr}^{v_{6}}\left(\neg v_{4}\right)$ for vertex $V_{4}$.
c. Now, suppose that after processing the evidence $V_{6}=$ true, the evidence $V_{2}=$ false is obtained. For each vertex in the network, examine whether or not this new evidence can influence the probabilities of its values.

## * Exercise 4.5

Consider the Bayesian network $\mathcal{B}=(G, \Gamma)$ where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is the following digraph:

and $\Gamma=\left\{\gamma_{V_{i}} \mid V_{i} \in \boldsymbol{V}_{G}\right\}$ is defined as

$$
\begin{array}{ll}
\gamma_{V_{1}}\left(v_{1}\right)=0.8 & \gamma_{V_{2}}\left(v_{2}\right)=0.5 \\
\gamma_{V_{4}}\left(v_{4}\right)=0.3 & \\
\gamma_{V_{3}}\left(v_{3} \mid v_{1} \wedge v_{2}\right)=0.2 & \gamma_{V_{6}}\left(v_{6} \mid v_{3} \wedge v_{4}\right)=0.3 \\
\gamma_{V_{3}}\left(v_{3} \mid \neg v_{1} \wedge v_{2}\right)=0.6 & \gamma_{V_{6}}\left(v_{6} \mid \neg v_{3} \wedge v_{4}\right)=0.8 \\
\gamma_{V_{3}}\left(v_{3} \mid v_{1} \wedge \neg v_{2}\right)=0.5 & \gamma_{V_{6}}\left(v_{6} \mid v_{3} \wedge \neg v_{4}\right)=0.2 \\
\gamma_{V_{3}}\left(v_{3} \mid \neg v_{1} \wedge \neg v_{2}\right)=0.1 & \gamma_{V_{6}}\left(v_{6} \mid \neg v_{3} \wedge \neg v_{4}\right)=0.1 \\
& \\
\gamma_{V_{5}}\left(v_{5} \mid v_{3}\right)=0.6 & \gamma_{V_{7}}\left(v_{7} \mid v_{5}\right)=0.9 \\
\gamma_{V_{5}}\left(v_{5} \mid \neg v_{3}\right)=0.4 & \gamma_{V_{7}}\left(v_{7} \mid \neg v_{5}\right)=0.8
\end{array}
$$

Let $\operatorname{Pr}$ be the joint probability distribution defined by the Bayesian network $\mathcal{B}$. Suppose that Pearl's algorithm is employed for probabilistic inference with $\mathcal{B}$.
a. Compute the probabilities $\operatorname{Pr}\left(v_{5}\right)$ and $\operatorname{Pr}\left(\neg v_{5}\right)$ for vertex $V_{5}$ from the Bayesian network $\mathcal{B}$.
b. Suppose that the evidence $V_{3}=$ true is observed and entered into the Bayesian network $\mathcal{B}$. Compute the updated probabilities $\operatorname{Pr}^{v_{3}}\left(v_{5}\right)$ and $\operatorname{Pr}^{v_{3}}\left(\neg v_{5}\right)$ for vertex $V_{5}$ from the network.
c. For each vertex in the network, examine whether or not the evidence $V_{3}=$ true may influence the probabilities of this vertex' values.

## * Exercise 4.6

Consider the Bayesian network $\mathcal{B}=(G, \Gamma)$ where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is the following digraph:

and $\Gamma=\left\{\gamma_{V_{i}} \mid V_{i} \in \boldsymbol{V}_{G}\right\}$ is defined as

$$
\begin{array}{ll}
\gamma_{V_{1}}\left(v_{1}\right)=0.5 & \gamma_{V_{5}}\left(v_{5} \mid v_{3}\right)=0.2 \\
& \gamma_{V_{5}}\left(v_{5} \mid \neg v_{3}\right)=1 \\
\gamma_{V_{2}}\left(v_{2} \mid v_{1}\right)=0.35 & \\
\gamma_{V_{2}}\left(v_{2} \mid \neg v_{1}\right)=0.1 & \gamma_{V_{6}}\left(v_{6} \mid v_{3} \wedge v_{4}\right)=0.75 \\
& \gamma_{V_{6}}\left(v_{6} \mid \neg v_{3} \wedge v_{4}\right)=0.3 \\
\gamma_{V_{3}}\left(v_{3} \mid v_{1}\right)=0.5 & \gamma_{V_{6}}\left(v_{6} \mid v_{3} \wedge \neg v_{4}\right)=0.25 \\
\gamma_{V_{3}}\left(v_{3} \mid \neg v_{1}\right)=0 & \gamma_{V_{6}}\left(v_{6} \mid \neg v_{3} \wedge \neg v_{4}\right)=0.2 \\
& \\
\gamma_{V_{4}}\left(v_{4}\right)=0.5 & \gamma_{V_{7}}\left(v_{7} \mid v_{6}\right)=0.8 \\
& \gamma_{V_{7}}\left(v_{7} \mid \neg v_{6}\right)=0
\end{array}
$$

Let $\operatorname{Pr}$ be the joint probability distribution defined by the Bayesian network $\mathcal{B}$. Suppose that Pearl's algorithm is employed for probabilistic inference with $\mathcal{B}$.
a. Suppose that the evidence $V_{7}=$ true has been entered into the network $\mathcal{B}$. Compute the probabilities $\operatorname{Pr}^{v_{7}}\left(v_{1}\right)$ and $\operatorname{Pr}^{v_{7}}\left(\neg v_{1}\right)$ for vertex $V_{1}$ from the Bayesian network $\mathcal{B}$.
b. Now suppose that the evidence $V_{4}=$ true is entered into the network as well. Does the probability $\operatorname{Pr}^{v_{7}, v_{4}}\left(v_{1}\right)$ differ from the probability $\operatorname{Pr}^{v_{7}}\left(v_{1}\right)$ ?
c. Suppose that, for the application for which the Bayesian network $\mathcal{B}$ has been developed, the probability $\operatorname{Pr}^{v_{7}, v_{4}}\left(v_{1} \wedge \neg v_{2}\right)$ needs to be computed. Explain how this probability can be computed from the network efficiently.

## * Exercise 4.7

Consider the Bayesian network $\mathcal{B}=(G, \Gamma)$ where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is the following digraph:

and $\Gamma=\left\{\gamma_{V_{i}} \mid V_{i} \in \boldsymbol{V}_{G}\right\}$ is defined as:

$$
\begin{array}{ll}
\gamma_{V_{1}}\left(v_{1}\right)=0.05 & \gamma_{V_{5}}\left(v_{5} \mid v_{3}\right)=0.35 \\
& \gamma_{V_{5}}\left(v_{5} \mid \neg v_{3}\right)=0 \\
\gamma_{V_{2}}\left(v_{2}\right)=0.3 & \\
\gamma_{V_{3}}\left(v_{3} \mid v_{1}\right)=0.75 & \gamma_{V_{6}}\left(v_{6} \mid v_{3}\right)=0.8 \\
\gamma_{V_{3}}\left(v_{3} \mid \neg v_{1}\right)=0.4 & \gamma_{V_{6}}\left(v_{6} \mid \neg v_{3}\right)=0.6 \\
& \\
\gamma_{V_{4}}\left(v_{4} \mid v_{1} \wedge v_{2}\right)=0.1 & \gamma_{V_{7}}\left(v_{7} \mid v_{6}\right)=0.5 \\
\gamma_{V_{4}}\left(v_{4} \mid \neg v_{1} \wedge v_{2}\right)=0.25 & \gamma_{V_{7}}\left(v_{7} \mid \neg v_{6}\right)=0.45 \\
\gamma_{V_{4}}\left(v_{4} \mid v_{1} \wedge \neg v_{2}\right)=0.8 & \\
\gamma_{V_{4}}\left(v_{4} \mid \neg v_{1} \wedge \neg v_{2}\right)=0.95 &
\end{array}
$$

Let $\operatorname{Pr}$ be the joint probability distribution defined by the Bayesian network B. Suppose that Pearl's algorithm is employed for probabilistic inference with $\mathcal{B}$.
a. Suppose that the evidence $V_{1}=$ true and $V_{6}=$ false has been entered into the Bayesian network $\mathcal{B}$. Compute the probabilities $\operatorname{Pr}^{v_{1}, \neg v_{6}}\left(v_{3}\right)$ and $\operatorname{Pr}^{v_{1}, \neg v_{6}}\left(\neg v_{3}\right)$ for vertex $V_{3}$ from $\mathcal{B}$.
b. Suppose that, for the application for which the Bayesian network $\mathcal{B}$ has been developed, the probability $\operatorname{Pr}^{v_{1}, \neg v_{6}}\left(v_{2} \wedge v_{5}\right)$ needs to be computed. Explain how this probability can be computed from the network efficiently.
c. Now suppose that, for the application for which the Bayesian network has been developed, the probability $\operatorname{Pr}^{v_{1}, \neg v_{6}}\left(v_{2} \vee v_{4}\right)$ needs to be computed. Explain how this probability can be computed from the network efficiently.

## * Exercise 4.8

Let $\mathcal{B}=(G, \Gamma)$ be a Bayesian network where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ with $\boldsymbol{V}_{G}=\left\{V_{1}, \ldots, V_{n}\right\}$, $n \geq 1$, and $\boldsymbol{A}_{G}=\left\{\left(V_{i}, V_{i+1}\right) \mid i=1, \ldots, n-1\right\}$ is the following digraph:


Let $\operatorname{Pr}$ be the joint probability distribution defined by the Bayesian network $\mathcal{B}$. Furthermore, let $\mathcal{B}^{\prime}=\left(G, \Gamma^{\prime}\right)$ be the Bayesian network that is obtained from $\mathcal{B}$ by substituting for a vertex $V_{k}, k>1$, the function value $\gamma_{V_{k}}\left(v_{k} \mid v_{k-1}\right)$ by $\gamma_{V_{k}}\left(v_{k} \mid v_{k-1}\right)+\epsilon$, for some small $\epsilon>0$. Let $\operatorname{Pr}^{\prime}$ be the joint probability distribution defined by $\mathcal{B}^{\prime}$. Show that $\left|\operatorname{Pr}^{\prime}\left(V_{i}\right)-\operatorname{Pr}\left(V_{i}\right)\right| \leq \epsilon$, for all $V_{i} \in \boldsymbol{V}_{G}, i=1, \ldots, n$.

## Exercise 4.9

Consider the Bayesian network $\mathcal{B}=(G, \Gamma)$ where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is the following digraph:

and $\Gamma=\left\{\gamma_{V_{i}} \mid V_{i} \in \boldsymbol{V}_{G}\right\}$ is defined as

$$
\begin{array}{ll}
\gamma_{V_{1}}\left(v_{1}\right)=0.5 \\
& \\
\gamma_{V_{2}}\left(v_{2} \mid v_{1}\right)=0.9 & \gamma_{V_{3}}\left(v_{3} \mid v_{2}\right)=0.5 \\
\gamma_{V_{2}}\left(v_{2} \mid \neg v_{1}\right)=0.7 & \gamma_{V_{3}}\left(v_{3} \mid \neg v_{2}\right)=0.8 \\
& \\
\gamma_{V_{4}}\left(v_{4} \mid v_{2} \wedge v_{3}\right)=0.1 & \gamma_{V_{5}}\left(v_{5} \mid v_{2} \wedge v_{3}\right)=0.3 \\
\gamma_{V_{4}}\left(v_{4} \mid \neg v_{2} \wedge v_{3}\right)=0.25 & \gamma_{V_{5}}\left(v_{5} \mid \neg v_{2} \wedge v_{3}\right)=0.7 \\
\gamma_{V_{4}}\left(v_{4} \mid v_{2} \wedge \neg v_{3}\right)=0.6 & \gamma_{V_{5}}\left(v_{5} \mid v_{2} \wedge \neg v_{3}\right)=0.2 \\
\gamma_{V_{4}}\left(v_{4} \mid \neg v_{2} \wedge \neg v_{3}\right)=0.75 & \gamma_{V_{5}}\left(v_{5} \mid \neg v_{2} \wedge \neg v_{3}\right)=0.7 \\
\gamma_{V_{6}}\left(v_{6} \mid v_{3}\right)=0.45 & \\
\gamma_{V_{6}}\left(v_{6} \mid \neg v_{3}\right)=0.65 &
\end{array}
$$

Let $\operatorname{Pr}$ be the joint probability distribution defined by the Bayesian network $\mathcal{B}$.

* a. Compute a loop cutset for the digraph $G$ by employing the heuristic algorithm of Suermondt and Cooper.
* b. Compute the probabilities $\operatorname{Pr}\left(v_{5}\right)$ and $\operatorname{Pr}\left(\neg v_{5}\right)$ for vertex $V_{5}$ from the Bayesian network $\mathcal{B}$ by Pearl's algorithm with loop cutset conditioning.
c. Suppose that the evidence $V_{3}=$ false is observed and entered into the network. Furthermore, suppose that after entering $V_{3}=$ false into the network, the additional evidence $V_{1}=$ true is observed. For each vertex in the network, examine whether or not this new evidence can influence the probabilities of its values.


## * Exercise 4.10

Consider the Bayesian network $\mathcal{B}=(G, \Gamma)$ where $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ is the following digraph:

and $\Gamma=\left\{\gamma_{V_{i}} \mid V_{i} \in \boldsymbol{V}_{G}\right\}$ is defined as

$$
\begin{array}{ll}
\gamma_{V_{1}}\left(v_{1}\right)=0.9 & \gamma_{V_{5}}\left(v_{5} \mid v_{2} \wedge v_{3} \wedge v_{4}\right)=0.2 \\
& \gamma_{V_{5}}\left(v_{5} \mid \neg v_{2} \wedge v_{3} \wedge v_{4}\right)=0.25 \\
\gamma_{V_{2}}\left(v_{2} \mid v_{1}\right)=0.125 & \gamma_{V_{5}}\left(v_{5} \mid v_{2} \wedge \neg v_{3} \wedge v_{4}\right)=0.7 \\
\gamma_{V_{2}}\left(v_{2} \mid \neg v_{1}\right)=0.875 & \gamma_{V_{5}}\left(v_{5} \mid v_{2} \wedge v_{3} \wedge \neg v_{4}\right)=0.6 \\
& \gamma_{V_{5}}\left(v_{5} \mid \neg v_{2} \wedge \neg v_{3} \wedge v_{4}\right)=0.55 \\
\gamma_{V_{3}}\left(v_{3} \mid v_{2}\right)=0.7 & \gamma_{V_{5}}\left(v_{5} \mid \neg v_{2} \wedge v_{3} \wedge \neg v_{4}\right)=0.8 \\
\gamma_{V_{3}}\left(v_{3} \mid \neg v_{2}\right)=0.2 & \gamma_{V_{5}}\left(v_{5} \mid v_{2} \wedge \neg v_{3} \wedge \neg v_{4}\right)=0.3 \\
& \gamma_{V_{5}}\left(v_{5} \mid \neg v_{2} \wedge \neg v_{3} \wedge \neg v_{4}\right)=0.15 \\
\gamma_{V_{4}}\left(v_{4} \mid v_{2}\right)=0.35 & \\
\gamma_{V_{4}}\left(v_{4} \mid \neg v_{2}\right)=0.1 &
\end{array}
$$

Let $\operatorname{Pr}$ be the joint probability distribution defined by the Bayesian network $\mathcal{B}$.
a. Compute a loop cutset for the digraph $G$ by employing the heuristic algorithm of Suermondt and Cooper.
b. Compute the probabilities $\operatorname{Pr}\left(v_{3}\right)$ and $\operatorname{Pr}\left(\neg v_{3}\right)$ for vertex $V_{3}$ from the Bayesian network $\mathcal{B}$ by Pearl's algorithm with loop cutset conditioning.
c. Suppose that the evidence $V_{5}=$ true is entered into the network. Explain how the probabilities $\operatorname{Pr}^{v_{5}}\left(v_{3}\right)$ and $\operatorname{Pr}^{v_{5}}\left(\neg v_{3}\right)$ are computed from the network.

## Exercise 4.11

* a. Consider a Bayesian network $\mathcal{B}=(G, \Gamma)$ where $G$ is the following digraph:


Compute a loop cutset for $G$ by employing the heuristic algorithm of Suermondt and Cooper.
b. The heuristic algorithm of Suermondt and Cooper aims at computing for a multiply connected digraph an optimal loop cutset. Now consider a Bayesian network in which the variables are not necessarily binary but may take a value from an arbitrarily large finite set of values. Illustrate by means of an example that for such a Bayesian network a non-optimal loop cutset may be a better choice for the purpose of loop cutset conditioning than an optimal loop cutset.
c. In the heuristic algorithm of Suermondt and Cooper the number of values for the various variables of a Bayesian network are not taken into consideration. Modify the algorithm in such a way that the number of values is taken into consideration to arrive at a 'good' loop cutset.

## Exercise 4.12

* a. Consider the following digraph $G$ :


Compute a loop cutset for $G$ by employing the heuristic algorithm of Suermondt and Cooper.
b. Let $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ be an acyclic multiply connected digraph and let $\boldsymbol{C} \subseteq \boldsymbol{V}_{G}$ be a set of vertices in $G$. Furthermore, let $G^{\prime}$ be the digraph that is obtained from $G$ by deleting for each vertex from $\boldsymbol{C}$ its outgoing arcs, that is, $G^{\prime}=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}^{\prime}\right)$ with $\boldsymbol{A}_{G}^{\prime}=\boldsymbol{A}_{G} \backslash\left\{\left(V_{i}, V_{j}\right) \mid V_{i} \in \boldsymbol{C}\right\}$. Show that the set $\boldsymbol{C}$ is a loop cutset of $G$ if and only if $G^{\prime}$ is a singly connected digraph.
c. Upon employing the heuristic algorithm of Suermondt and Cooper for computing a loop cutset for a multiply connected digraph $G$, a non-minimal loop cutset $\boldsymbol{L}_{G}$ may be yielded. Give a sketch of an algorithm that takes for its input the non-minimal loop cutset $\boldsymbol{L}_{G}$ and yields a minimal loop cutset $\boldsymbol{L}_{G}^{\prime}$ of $G$.

## Exercise 4.13

For probabilistic inference with a Bayesian network, sometimes Pearl's algorithm is supplemented with a method called evidence absorption. The basic idea of this method is to dynamically modify a Bayesian network after evidence has been entered so as to reflect as many of the new independences occasioned by the evidence as possible. Now consider a Bayesian network $\mathcal{B}=(G, \Gamma)$ and suppose that the evidence $V_{i}=$ true is entered for some vertex $V_{i} \in \boldsymbol{A}_{G}$. Evidence absorption now amounts to modifying the digraph $G$ of $\mathcal{B}$ by deleting all outgoing arcs from vertex $V_{i}$; with the thus modified digraph $G^{\prime}$ a new set $\Gamma^{\prime}$ of probability assessment functions is constructed such that the new Bayesian network $\mathcal{B}^{\prime}=\left(G^{\prime}, \Gamma^{\prime}\right)$ reflects the updated joint probability distribution.
a. Show that the digraphs $G$ and $G^{\prime}$ portray the same set of independences given $V_{i}$.

Hint. Show that: $\langle\boldsymbol{X}| \boldsymbol{Y}|\boldsymbol{Z}\rangle_{G}^{d}$ if and only if $\langle\boldsymbol{X}| \boldsymbol{Y}|\boldsymbol{Z}\rangle_{G^{\prime}}^{d}$, for all sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z} \subseteq \boldsymbol{V}_{G}$ with $V_{i} \in \boldsymbol{Y}$.
b. Show by means of an example that the digraphs $G$ and $G^{\prime}$ do not represent the same set of independences.
c. Explain why repeated application of the method of evidence absorption reduces the average-case computational complexity of reasoning with a Bayesian network.

## Chapter 5

## Building a Bayesian Network

The Bayesian network framework generally is used for applying probability theory for reasoning with uncertainty in knowledge-based systems. For employing the framework for a real-life domain of application, relevant knowledge of the domain at hand is represented in the Bayesian network formalism; the basic algorithms of the framework are taken as building blocks for shaping the system's intelligent problem-solving behaviour. In this chapter we address the task of building a Bayesian network for a domain of application; shaping intelligent problem-solving behaviour is one of the topics of the next chapter.

Building a Bayesian network for a domain of application involves various different tasks. The first of these tasks is to identify the variables that are of importance in the domain at hand, along with their possible values. Once the important domain variables have been identified, the qualitative part of the Bayesian network in the making is constructed: the second task in building a Bayesian network therefore is to identify the independences among the variables discerned and to express these in an acyclic digraph. After the qualitative part of the network has been established, the probability assessment functions are defined: the last task in building a Bayesian network is to estimate the (conditional) probabilities that are required to constitute the network's quantitative part. The three tasks are summarised in Figure 5.1. The various tasks in building a Bayesian network are, in principle, performed one after the other. Building a network, however, often requires a careful trade-off between the desire for a large and rich model to obtain accurate results on the one hand, and the costs of construction and maintenance, and the complexity of probabilistic inference on the other hand. In practice, therefore, building a Bayesian network is a cyclic process that iterates over these tasks until a network results that is deemed satisfactory for the domain at hand. We would like to note that, although the Bayesian network framework has been around for some time, methodologies for building real-life Bayesian networks do not yet abound. Building a Bayesian network is an engineering task and it therefore makes sense to adopt common practices from(software) engineering, especially for building large models [Laskey \& Mahoney, 2000]. Recently, various Bayesian networks have been built for environmental applications. This led to the publication of a paper with guidelines to developing and evaluating Bayesian network models of environmental systems [Chen \& Pollino, 2012]; the paper summarises most of the steps detailed in this section and includes lots of useful references.

In the following sections, we address the tasks involved in building a Bayesian network separately. In Section 5.1 we address the identification of the important domain


Figure 5.1: The tasks in building a Bayesian network.
variables and values. In Section 5.2 we discuss the construction of the qualitative part of a Bayesian network in the making. Section 5.3 focuses on the assessment of the probabilities required for the network's quantitative part.

### 5.1 Identifying Variables and Values

The first task in building a Bayesian network for a domain of application is to identify the variables that are of importance in the domain, along with their possible values. Identifying the important domain variables is typically performed with the help of one or more experts. Since a Bayesian network, as any model, necessarily is a simplification of reality, well-founded decisions have to be taken as to which variables should be included in the network and which may be omitted. We would like to note that this task is not reserved for building Bayesian networks, but instead is quite common in engineering knowledge-based systems. A knowledge engineer can therefore make use of the various elicitation techniques designed for engineering knowledge-based systems in general. For example, the use of ontologies was suggested to help structure the domain, its variables and the different types of relations that exist between the variables [Helsper \& Van der Gaag, 2002]. Such ontologies have proved useful tools upon constructing a Bayesian network's digraph, for distilling the variables and their interrelationships that are important for the problem at hand. We will not further elaborate on this task here and confine ourselves to emphasising that it needs to be performed with care since the domain variables that are modelled in the Bayesian network demarcate the scope of the resulting system.

Once the important domain variables have been identified, each of them has to be expressed as a statistical variable to allow for inclusion in the Bayesian network in the making. We recall that a statistical variable is characterised by its values being mutually exclusive and collectively exhaustive. To allow for inclusion in a Bayesian network, a statistical variable has to take its value from a finite set of discrete values. Only if the set of values of a domain variable satisfies these properties, can the variable
be included in the network as it is. Otherwise, the domain variable is modelled in terms of one or more statistical variables.

For modelling domain variables for inclusion in a Bayesian network, we distinguish between the following types of variable:

- a single-valued domain variable that takes a value from a finite set of (discrete) values;
- a single-valued domain variable that takes a value from an infinite set of values;
- a multi-valued domain variable.

We begin by considering a single-valued domain variable that takes its value from a finite set of discrete values. The values of this variable are mutually exclusive by definition. Moreover, the variable's values are easily rendered collectively exhaustive. A single-valued domain variable can therefore be expressed straightforwardly as a statistical variable in the Bayesian network in the making.

For a single-valued domain variable that takes its value from an infinite set of values, we equally have that this variable's values are mutually exclusive and are easily rendered collectively exhaustive. However, its set of possible values being infinite prohibits including the domain variable as it is in the Bayesian network in the making. To allow for inclusion in the network, the set of values is discretised. To this end, the variable's possible values are clustered in a finite number of mutually exclusive and collectively exhaustive clusters of values; these clusters then are taken as the values of a new single-valued variable that is included in the Bayesian network in the making.

Example 5.1.1 Suppose that, in a medical domain of application, a patient's temperature has been identified as an important domain variable. This domain variable is single-valued, taking its value from the (in essence) infinite range of values between, say, $35^{\circ} \mathrm{C}$ and $41^{\circ} \mathrm{C}$. This range of values can be discretised in, for example, the five intervals $[35 ; 36),[36 ; 37),[37 ; 38),[38 ; 40)$, and $[40 ; 41]$. Note that these intervals can be looked upon as values that are mutually exclusive and collectively exhaustive. For inclusion in the Bayesian network in the making, the domain variable is now expressed as a statistical variable that takes its value from among the five values mentioned above.

In general, a domain variable's set of values admits numerous discretisations. The discretisation best chosen is highly dependent on the purpose for which the Bayesian network is being developed and has to be decided upon in consultation with one or more domain experts.

A domain variable taking multiple values from a finite set of discrete values cannot be expressed directly as a statistical variable since its values are not mutually exclusive. To allow for inclusion in the Bayesian network in the making, such a domain variable needs to be re-modelled so as to satisfy the properties of a statistical variables. Remodelling may be done by redefining the variable's set of values in such a way that the new values are mutually exclusive and collectively exhaustive.

Example 5.1.2 Suppose that, in a medical domain of application, a patient's blood count has been identified as an important domain variable. This domain variable is multi-valued, taking its values from among the set of values \{normal, lymphocytosis,
lymphocytopenia, leucocytosis, leucocytopenia\}. The values lymphocytosis and lymphocytopenia refer to the level of lymphocytes in a patient's blood sample and the values leucocytosis and leucocytopenia refer to the leucocyte level; these values are not mutually exclusive. Domain knowledge now indicates that a patient's blood count cannot take any arbitrary combination of the values mentioned above. In fact, only nine combinations of values are meaningful - these are \{normal\}, \{lymphocytosis\}, $\{$ lymphocytopenia\}, $\{$ leucocytosis $\},\{$ leucocytopenia $\},\{$ lymphocytosis, leucocytosis $\}$, \{lymphocytosis, leucocytopenia\}, \{lymphocytopenia, leucocytosis\}, and \{lymphocytopenia, leucocytopenia\}. These combinations are now taken as the mutually exclusive values of a new statistical variable, modelling a patient's blood count, to be included in the Bayesian network in the making.

Multi-valuedness of a domain variable often arises from the variable being compound, that is, implicitly built from several other domain variables. Instead of redefining its set of values, a compound variable may be modelled by decomposing it into the domain variables it is built from and subsequently modelling these separate variables in the Bayesian network in the making.

Example 5.1.3 We consider once more the domain variable expressing a patient's blood count as introduced in the previous example. We observe that this domain variable is a compound variable built from two variables. One of these variables expresses the level of lymphocytes in a patient's blood sample, taking its value from the set of values \{normal, lymphocytosis, lymphocytopenia\}; the other variable captures the level of leucocytes, taking its value from the set of values \{normal, leucocytosis, leucocytopenia\}. Note that these two variables are statistical variables, that can be included directly in the Bayesian network in the making. The originally multi-valued domain variable is now decomposed into these variables.

Decomposing a compound variable requires considerable knowledge of the domain at hand and therefore is best performed with the help of a domain expert.

After all important domain variables have been expressed as statistical variables for inclusion in the Bayesian network in the making, the meanings of these statistical variables have to be properly documented. While each statistical variable has a unique meaning in the context of the Bayesian network at hand, its meaning may very well differ from the meanings associated with variables with the same name in other models or in the literature in the domain. The task of documenting the meanings of the statistical variables modelled in the Bayesian network therefore is crucial to avoid any ambiguity in future reference. This task is not typical for building Bayesian networks and has in fact been recognised as vital in knowledge-engineering practice in general.

### 5.2 Constructing the Digraph

Once the domain variables of importance have been identified and expressed as statistical variables, the construction of the qualitative part of the Bayesian network in the making can commence. In principle, for constructing the network's qualitative part, the independence relation of the joint probability distribution on the variables discerned has to be identified and represented in an acyclic digraph. In practice, however, the digraph typically is conceived without explicitly identifying independences: it is constructed directly, either by hand or by learning from data.

### 5.2.1 Constructing the Digraph by Hand

For most domains of application, the qualitative part of a Bayesian network in the making is handcrafted. For this purpose, one or more domain experts are interviewed. In the interviews, the qualitative relationships between the variables discerned are elicited, using the concept of causality as a heuristic guiding principle. Typical questions asked during the interviews are "What could cause this effect ?", "What manifestations could this cause have ?" [Henrion, 1989]. The elicited causal relationships among the variables are expressed in graphical terms by taking the direction of causality for directing the arcs between related variables. The resulting graphical representation can then be taken as a basis for further refinement.

Using the concept of causality as a guiding principle during the interviews with domain experts may meet with some difficulties. For example, not every influential relationship among variables can be interpreted as causal. If a non-causal influential relationship comes to the fore during an interview, a more elaborate analysis of the independences involved is required before it can be expressed in graphical terms. A tool that can support this analysis by visually and verbally explaining the (in)dependences captured in the graphical structure is Matilda [Boneh, Nicholson \& Sonenberg, 2006]. Another problem that may arise from building on causality concerns the creation of cycles in the graphical representation. As the digraph of a Bayesian network has to be acyclic, any cycle that has resulted from the interviews needs to be removed. There are several ways to deal with cycles [Peek \& Ottenkamp, 1997], the simplest being the deletion of a single arc from a cycle to be removed. To conclude, causality is not a well-understood concept and therefore may leave room for multiple interpretations.

Despite the fact that building on the concept of causality does not suffice in itself for constructing the digraph of a Bayesian network in the making, it brings the advantage that domain experts are allowed to express their knowledge in either the causal or diagnostic direction. Since they are allowed to express their knowledge in a form they feel comfortable with, the experts' statements tend to be quite robust. We would like to note that the task of eliciting relationships among variables from domain experts is not reserved for building Bayesian networks: the elicitation task has parallels with engineering knowledge-based systems in general for which several methodologies have been developed [Boose \& Gaines, 1988]. Some experiences with the issues discussed above in building a Bayesian network for a real-life application are described in [Van der Gaag \& Helsper, 2002].

### 5.2.2 Learning the Digraph from Data

In various domains of application, data has been collected and maintained over numerous years of every-day problem solving. Such a data collection usually contains highly valuable information about the relationships among the variables discerned, be it implicitly. If a comprehensive data set is available in the domain for which a Bayesian network is developed, the construction of the network's qualitative part may be performed automatically: the basic idea of learning the qualitative part from data is to distil information from the data set and exploit it for constructing a digraph.

## The data set

For learning the qualitative part of a Bayesian network, a data set with information from the domain at hand is required.

We formally define the concept of a database.
Definition 5.2.1 Let $\boldsymbol{V}$ be a set of variables. A data set $\boldsymbol{D}$ over $\boldsymbol{V}$ is a multi-set of configurations of $\boldsymbol{V}$. An element $c_{V}$ of $\boldsymbol{D}$ is called a case.

To be suitable for learning purposes, a data set has to satisfy various properties. First of all, the data comprised in the data set must have been collected very carefully. Biases that are introduced in the data set as a result of the data collection strategies used will have impact on the topology of the resulting digraph, yet may not be desirable for the purpose for which the Bayesian network is being developed; unfortunately, biases are not easily detected in a once constructed network. Also, to be suitable for the purpose of learning a Bayesian network's digraph, the variables and associated values that occur in the data set should match the variables and values that are to be modelled in the network, or should at least admit easy translation into these variables and values. Moreover, the data set should comprise enough data to allow for reliable identification of probabilistic relationships among the variables discerned.

In addition to the general prerequisites outlined before, a data set should satisfy several properties that are implicitly assumed by most algorithms for learning a Bayesian network's digraph from data. A property that is commonly assumed is that each case in the data set specifies a value for every variable discerned, that is, there are no missing values. Unfortunately, for most real-life data sets this property does not hold. To use a data set with missing values for learning purposes, therefore, the missing values have to be filled in, for example based upon (roughly) estimated prior probabilities for these values or with the help of domain experts; the EM (Expectation-Maximisation) algorithm is often used for this purpose [Dempster et al., 1977]. Other properties assumed by most learning algorithms concern the real-life process that generated the cases comprised in the data set at hand. It is assumed that this process generates cases independently, that is, the values specified for the variables in a case are assumed not to be influenced in any way by the values in previously generated cases. Also, it is assumed that the process is not time-dependent, that is, the data set is assumed not to reflect information that varies over time.

As we will discuss in the sequel, a data set of cases over a set of statistical variables is used for estimating various conditional probabilities. These conditional probabilities are used for analysing the strengths of various different relationships among the variables. The conditional probabilities required for this purpose are obtained from the data set by counting. We consider a data set $\boldsymbol{D}$ over the set of variables $\boldsymbol{V}$, comprising $N$ cases. For each set of variables $\boldsymbol{X} \subseteq \boldsymbol{V}$, we will use $N\left(c_{\boldsymbol{X}}\right)$ to denote the number of cases in $\boldsymbol{D}$ in which the variables from $\boldsymbol{X}$ have adopted the conjunction of values $c_{\boldsymbol{X}}$; for the empty set, we take $N\left(c_{\varnothing}\right)=N$. For any two sets of variables $\boldsymbol{X}, \boldsymbol{Y} \subseteq \boldsymbol{V}$, the probability distribution $\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Y})$ can now be estimated from the data set by

$$
\operatorname{Pr}\left(c_{\boldsymbol{X}} \mid c_{\boldsymbol{Y}}\right)=\frac{N\left(c_{\boldsymbol{X}} \wedge c_{\boldsymbol{Y}}\right)}{N\left(c_{\boldsymbol{Y}}\right)}
$$

for all configurations $c_{\boldsymbol{X}}$ and $c_{\boldsymbol{Y}}$ of $\boldsymbol{X}$ and $\boldsymbol{Y}$, respectively.
Various algorithms exist for learning Bayesian networks from data. Roughly, two approaches can be distinguished: one where statistical tests are applied to the data
to determine the independences; the other where the space of possible networks is searched for a good scoring network. Research into such learning algorithms is very popular now that so many data sets are available. Here, we will focus on a score \& search-based method that lies at the basis of most algorithms that are popular today (see for example [Stahlschmidt et al., 2013] for a description of various algorithms).

## The quality measure

An algorithm for learning the qualitative part of a Bayesian network from data typically generates various different acyclic digraphs and compares these as to their ability to describe or explain the data at hand. We will focus on the task of comparing digraphs before turning to their generation.

For comparing digraphs as to their ability to describe the data from the database at hand, a quality measure is employed. A quality measure is a function that assigns to a digraph a numerical value expressing how well this digraph fits the data; this numerical value is called the quality of the digraph given the database. In the literature, various different quality measures have been proposed, originating from different theories. Here, we will discuss only the $M D L$ quality measure. This measure originates from coding theory and is built on the minimum description length principle. For an overview of the most popular quality measures in use for learning the qualitative part of a Bayesian network from data, we refer the reader to [Bouckaert, 1995].

We define the MDL quality measure to pertain to an acyclic digraph and a database.
Definition 5.2.2 Let $\boldsymbol{V}=\left\{V_{1}, \ldots, V_{n}\right\}$, $n \geq 1$, be a set of statistical variables. Let $\boldsymbol{D}$ be a data set over $\boldsymbol{V}$ and let $N$ be the number of cases in $\boldsymbol{D}$. Furthermore, let $P$ be a probability distribution over the set of acyclic digraphs with vertex set $\boldsymbol{V}$. Let $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ be an acyclic digraph with $\boldsymbol{V}_{G}=\boldsymbol{V}$. Then, the quality of $G$ given $\boldsymbol{D}$, denoted $Q(G, \boldsymbol{D})$, is defined by

$$
Q(G, \boldsymbol{D})=\log P(G)-N \cdot H(G, \boldsymbol{D})-\frac{1}{2} K \cdot \log N
$$

where

$$
H(G, \boldsymbol{D})=-\sum_{V_{i} \in \boldsymbol{V}} \sum_{c_{V_{i}}} \sum_{c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}}\left(\frac{N\left(c_{V_{i}} \wedge c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)}{N}\right) \cdot \log \left(\frac{N\left(c_{V_{i}} \wedge c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)}{N\left(c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)}\right)
$$

and

$$
K=\sum_{V_{i} \in \boldsymbol{V}} 2^{\left|\boldsymbol{\rho}_{G}\left(V_{i}\right)\right|}
$$

From the previous definition, we have that the quality $Q(G, \boldsymbol{D})$ of an acyclic digraph $G$ given a data set $\boldsymbol{D}$ involves three terms: $\log P(G),-N \cdot H(G, \boldsymbol{D})$ and $-\frac{1}{2} K \cdot \log N$. We will take a closer look at these terms separately.

In the first term, $\log P(G)$, of the quality of a digraph $G$ given a data set $\boldsymbol{D}$, the probability $P(G)$ denotes the prior probability of the digraph $G$ being the qualitative part of a Bayesian network that describes the joint probability distribution that generated the data set $\boldsymbol{D}$. This term therefore expresses information about the 'real' qualitative part prior to observation of the database. In a real-life application of a learning algorithm, this term is used for modelling domain knowledge. For example,
if a domain expert suggests existence of a specific arc in the qualitative part of the 'real' network, then digraphs that adhere to this suggestion are given a higher prior probability than digraphs that do not. In the sequel, we will assume that no prior domain knowledge is available and, hence, that the probability distribution $P$ is a uniform distribution.

The second term, $-N \cdot H(G, \boldsymbol{D})$, of the quality of a digraph $G$ given a data set $\boldsymbol{D}$ is proportional to the probability of the data set $\boldsymbol{D}$ being generated by the joint probability distribution represented by a Bayesian network $\mathcal{B}$ that involves the digraph $G$. Now, recall that a Bayesian network not only involves a digraph but also includes a set of probability assessment functions. To enhance the network's 'match' with the data set, all conditional probabilities required are assessed from the data: for each variable $V_{i} \in \boldsymbol{V}$, we take

$$
\gamma_{V_{i}}\left(c_{V_{i}} \mid c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)=\frac{N\left(c_{V_{i}} \wedge c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)}{N\left(c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)}
$$

for all configurations $c_{V_{i}}$ of $V_{i}$ and all configurations $c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}$ of the set $\boldsymbol{\rho}_{G}\left(V_{i}\right)$ of $V_{i}$ 's predecessors in $G$. Now consider the probability of the data set $\boldsymbol{D}$ being generated by the joint probability distribution $\operatorname{Pr}$ defined by the Bayesian network $\mathcal{B}=(G, \Gamma)$ where $\Gamma$ is the set of probability assessment functions estimated from $\boldsymbol{D}$. Since we have assumed that all cases in $\boldsymbol{D}$ have been generated independently, we have that the probability $P^{\prime}(\boldsymbol{D} \mid \mathcal{B})$ of the data set $\boldsymbol{D}$ given the Bayesian network $\mathcal{B}$ equals

$$
P^{\prime}(\boldsymbol{D} \mid \mathcal{B})=\prod_{c_{\boldsymbol{V}} \in \boldsymbol{D}} \operatorname{Pr}\left(c_{\boldsymbol{V}}\right)
$$

By exploiting the property stated in Proposition 4.1.3, we find that

$$
P^{\prime}(\boldsymbol{D} \mid \mathcal{B})=\prod_{c_{\boldsymbol{V}} \in \boldsymbol{D}} \prod_{V_{i} \in \boldsymbol{V}} \gamma_{V_{i}}\left(c_{V_{i}} \mid c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)
$$

where $\bigwedge_{V_{i} \in \boldsymbol{V}}\left(c_{V_{i}} \wedge c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right) \equiv c_{\boldsymbol{V}}$. By reordering terms, we ultimately find that

$$
P^{\prime}(\boldsymbol{D} \mid \mathcal{B})=\prod_{V_{i} \in \boldsymbol{V}} \prod_{c_{V_{i}}} \prod_{c_{\boldsymbol{\rho}_{G}}\left(V_{i}\right)} \gamma_{V_{i}}\left(c_{V_{i}} \mid c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)^{N\left(c_{V_{i}} \wedge c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)}
$$

Substitution of the probability estimates for the function values of the various different assessment functions now yields

$$
P^{\prime}(\boldsymbol{D} \mid \mathcal{B})=\prod_{V_{i} \in \boldsymbol{V}} \prod_{c_{V_{i}}} \prod_{\boldsymbol{\rho}_{\boldsymbol{G}}\left(V_{i}\right)}\left(\frac{N\left(c_{V_{i}} \wedge c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)}{N\left(c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)}\right)^{N\left(c_{V_{i}} \wedge c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)}
$$

Hence,

$$
\log P^{\prime}(\boldsymbol{D} \mid \mathcal{B})=N \cdot \sum_{V_{i} \in \boldsymbol{V}} \sum_{c_{V_{i}}} \sum_{c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}}\left(\frac{N\left(c_{V_{i}} \wedge c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)}{N}\right) \cdot \log \left(\frac{N\left(c_{V_{i}} \wedge c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)}{N\left(c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)}\right)
$$

The result equals

$$
\log P^{\prime}(\boldsymbol{D} \mid \mathcal{B})=-N \cdot H(G, \boldsymbol{D})
$$

where $H(G, \boldsymbol{D})$ denotes the mutual entropy of $G$ and $\boldsymbol{D}$.
Entropy is an information-theoretic measure of uncertainty. In general, the entropy of a variable is maximal when the uncertainty concerning its value is maximal; the entropy is zero when there is complete knowledge as to the variable's value. Informally speaking, we have, in the context of the MDL quality measure, that the better a digraph fits the database, the lower their mutual entropy. In general, an acyclic digraph with more arcs will have a lower entropy than an acyclic digraph with fewer arcs: a digraph with more arcs expresses fewer independences and therefore is capable of capturing more of the nuances reflected in the cases of the data set. So, for a tightly connected acyclic digraph $G$ the term $-N \cdot H(G, \boldsymbol{D})$ in this digraph's quality tends to approach zero.

For an acyclic digraph with more arcs more conditional probabilities have to be estimated from the data set to define the probability assessment functions to be associated with the digraph, than for a digraph with fewer arcs. The more probabilities have to be estimated from the data set, the smaller the number of cases in the data set these estimates can be based upon. So, the more probabilities have to be estimated the less reliable the estimates will be. The third term, $-\frac{1}{2} K \cdot \log N$, of the quality $Q(G, \boldsymbol{D})$ of a digraph $G$ given a database $\boldsymbol{D}$ captures this aspect. In this term, $K$ expresses the number of probabilities that have to be assessed from the data set to define the probability assessment functions for the digraph at hand. So, the more arcs a digraph has, the larger $K$ and the smaller the term $-\frac{1}{2} K \cdot \log N$. The term $-\frac{1}{2} K \cdot \log N$ is often referred to as the penalty term as it imposes a penalty on adding arcs to the digraph in the making.

In the sequel, we shall see that an algorithm for learning the digraph of a Bayesian network from data constructs a digraph by successively adding single arcs to an initially arcless graph. For the initial arcless digraph, the entropy term $-N \cdot H(G, \boldsymbol{D})$ of its quality will be extremely small and the term $-\frac{1}{2} K \cdot \log N$ will be quite close to zero. On successively adding arcs to the digraph, the term $-N \cdot H(G, \boldsymbol{D})$ will grow rapidly. The term $-\frac{1}{2} K \cdot \log N$ on the other hand will decrease as the digraph becomes more tightly connected. As long as the entropy term increases more rapidly than the penalty term decreases, the quality of the digraph in the making will increase. The increase of the digraph's quality upon successive arc addition continues until the increase of the entropy term is dominated by the decrease of the penalty term.

## The search heuristic

We have mentioned before that an algorithm for learning the qualitative part of a Bayesian network from data generates various different acyclic digraphs and compares these digraphs with respect to their quality given the database at hand. The basic purpose of the algorithm is to select from among all possible acyclic digraphs a digraph with highest quality. Unfortunately, it is not feasible from a computational point of view to generate all digraphs and compute their qualities. A learning algorithm therefore incorporates a search heuristic that searches the set of all possible acyclic digraphs for digraphs that are likely to have a high quality; only for these digraphs is the quality given the data set actually computed. Several such search heuristics have been proposed in the literature. Here, we will only discuss the $B$ search heuristic. For an overview of various search heuristics and their properties, we refer the reader once more to [Bouckaert, 1995].

Let $\boldsymbol{V}$ be the set of statistical variables of the Bayesian network in the making. The B search heuristic for selecting a digraph of high quality from among all possible acyclic digraphs with vertex set $\boldsymbol{V}$ commences with investigating the arcless digraph $G_{0}=(\boldsymbol{V}, \varnothing)$. To this initially arcless graph, the search heuristic successively adds single arcs to improve the digraph's quality. Now, suppose that after adding some arcs, a digraph $G_{k}=\left(\boldsymbol{V}, \boldsymbol{A}_{G_{k}}\right)$ has resulted. For selecting a new arc to add to the digraph, the search heuristic identifies all possible arcs that can be added to the digraph $G_{k}$ without introducing a cycle. For each identified arc $\left(V_{i}, V_{j}\right)$, it computes the gain in quality that would be yielded by adding this arc to the digraph in the making. So, it computes the difference in quality between the new digraph $G_{k+1}=\left(\boldsymbol{V}, \boldsymbol{A}_{G_{k}} \cup\left\{\left(V_{i}, V_{j}\right)\right\}\right)$ and the digraph $G_{k}$. The search heuristic then selects an arc that yields the highest gain in quality and adds this arc to the digraph in the making. This process of arc addition is repeated until no gain in quality can be achieved any more. Note that the B search heuristic is a greedy search heuristic in the sense that it considers single arcs only.

From the informal outline of the B search heuristic, it will be evident that for selecting a digraph with high quality, the search heuristic has to compare qualities of various different digraphs. We observe, however, that for selecting the next arc to add to the digraph in the making the search heuristic does not actually need to know the qualities of the various digraphs considered since an arc is selected on the basis of differences in quality between digraphs only. From B being a greedy search heuristic, we further have that these differences are only computed for pairs of digraphs that differ in one arc. The B search heuristic exploits these observations by only partially computing the qualities of the digraphs concerned.

For the purpose of partially computing qualities, the B search heuristic makes use of the concept of the quality of a vertex in a digraph.

Definition 5.2.3 Let $\boldsymbol{V}=\left\{V_{1}, \ldots, V_{n}\right\}, n \geq 1$, be a set of statistical variables. Let $\boldsymbol{D}$ be a data set over $\boldsymbol{V}$ and let $N$ be the number of cases in $\boldsymbol{D}$. Let $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ be an acyclic digraph with $\boldsymbol{V}_{G}=\boldsymbol{V}$. For each vertex $V_{i} \in \boldsymbol{V}_{G}$, the quality of $V_{i}$ in $G$ given $\boldsymbol{D}$, denoted by $q\left(V_{i}, \boldsymbol{\rho}_{G}\left(V_{i}\right), \boldsymbol{D}\right)$, is defined as

$$
\begin{aligned}
q\left(V_{i}, \boldsymbol{\rho}_{G}\left(V_{i}\right), \boldsymbol{D}\right)=\sum_{c_{V_{i}}} \sum_{c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}} & N\left(c_{V_{i}} \wedge c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right) \cdot \log \left(\frac{N\left(c_{V_{i}} \wedge c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)}{N\left(c_{\boldsymbol{\rho}_{G}}\left(V_{i}\right)\right)}\right)+ \\
& -\frac{1}{2} \cdot 2^{\left|\boldsymbol{\rho}_{G}\left(V_{i}\right)\right|} \cdot \log N
\end{aligned}
$$

The quality of an acyclic digraph given a data set can now be expressed in terms of the qualities of its various vertices; this property is known as the sum property of a digraph's quality given a database and is stated more formally in the following lemma.

Lemma 5.2.4 Let $\boldsymbol{V}=\left\{V_{1}, \ldots, V_{n}\right\}, n \geq 1$, be a set of statistical variables. Let $\boldsymbol{D}$ be a data set over $\boldsymbol{V}$ and let $N$ be the number of cases in $\boldsymbol{D}$. Let $P$ be a probability distribution over the set of all acyclic digraphs with vertex set $\boldsymbol{V}$. Let $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ be such an acyclic digraph with $\boldsymbol{V}_{G}=\boldsymbol{V}$. Furthermore, let $Q(G, \boldsymbol{D})$ be the quality of $G$ given $\boldsymbol{D}$ and, for each vertex $V_{i} \in \boldsymbol{V}_{G}$, let $q\left(V_{i}, \boldsymbol{\rho}_{G}\left(V_{i}\right), \boldsymbol{D}\right)$ be the quality of $V_{i}$ in $G$ given $\boldsymbol{D}$. Then,

$$
Q(G, \boldsymbol{D})=\log P(G)+\sum_{V_{i} \in \boldsymbol{V}} q\left(V_{i}, \boldsymbol{\rho}_{G}\left(V_{i}\right), \boldsymbol{D}\right)
$$

Proof. The property stated in the lemma follows directly from Definition 5.2.2 and Definition 5.2.3.

Note that since we have assumed that the probability distribution $P$ over the set of all acyclic digraphs with a given vertex set is a uniform distribution, the term $\log P(G)$ in a digraph's quality is the same for every digraph $G$ and therefore is a constant with respect to learning the digraph of a Bayesian network from a data set.

Let $\boldsymbol{V}$ once more be the set of statistical variables discerned and let $G_{k}=\left(\boldsymbol{V}, \boldsymbol{A}_{G_{k}}\right)$ be a digraph that has resulted at some stage during application of the B search algorithm. Now, consider a new digraph $G_{k+1}=\left(\boldsymbol{V}, \boldsymbol{A}_{G_{k}} \cup\left\{\left(V_{i}, V_{j}\right)\right\}\right)$ that differs from $G_{k}$ in the $\operatorname{arc}\left(V_{i}, V_{j}\right)$ only. From Lemma 5.2 .4 it is easily seen that the difference in quality between these two digraphs equals the difference in quality of vertex $V_{j}$ in the two digraphs only; this property is stated in the following corollary.

Corollary 5.2.5 Let $\boldsymbol{V}=\left\{V_{1}, \ldots, V_{n}\right\}, n \geq 1$, be a set of statistical variables. Let $\boldsymbol{D}$ be a database over $\boldsymbol{V}$. Let $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ and $G^{\prime}=\left(\boldsymbol{V}_{G^{\prime}}, \boldsymbol{A}_{G} \cup\left\{\left(V_{i}, V_{j}\right)\right\}\right)$ be acyclic digraphs with $\boldsymbol{V}_{G}=\boldsymbol{V}_{G^{\prime}}=\boldsymbol{V}$ Furthermore, let $Q(G, \boldsymbol{D})$ and $Q\left(G^{\prime}, \boldsymbol{D}\right)$ be the quality of $G$ given $\boldsymbol{D}$ and the quality of $G^{\prime}$ given $\boldsymbol{D}$, respectively. For each $V_{k} \in \boldsymbol{V}$, let $q\left(V_{k}, \boldsymbol{\rho}_{G}\left(V_{k}\right), \boldsymbol{D}\right)$ and $q\left(V_{k}, \boldsymbol{\rho}_{G^{\prime}}\left(V_{k}\right), \boldsymbol{D}\right)$ be the quality of vertex $V_{k}$ in $G$ and the quality of $V_{k}$ in $G^{\prime}$ given $\boldsymbol{D}$, respectively. Then,

$$
Q\left(G^{\prime}, \boldsymbol{D}\right)-Q(G, \boldsymbol{D})=q\left(V_{j}, \boldsymbol{\rho}_{G^{\prime}}\left(V_{j}\right), \boldsymbol{D}\right)-q\left(V_{j}, \boldsymbol{\rho}_{G}\left(V_{j}\right), \boldsymbol{D}\right)
$$

The B search heuristic now exploits the property stated in Corollary 5.2.5 for computing the gain in quality that would result from adding a new arc to the digraph in the making.

The following pseudocode summarises the B search heuristic.

```
procedure construct-digraph \((\boldsymbol{V}, \boldsymbol{D}, G)\)
    for each variable \(V_{i} \in \boldsymbol{V}\) do
        \(\boldsymbol{\rho}_{G}\left(V_{i}\right):=\varnothing\)
    od;
    repeat
        for each (ordered) pair \(V_{i}, V_{j} \in \boldsymbol{V}\) such that addition of the \(\operatorname{arc}\left(V_{i}, V_{j}\right)\)
                to \(G\) does not introduce a (directed) cycle do
                \(\operatorname{diff}\left(V_{i}, V_{j}\right):=q\left(V_{j}, \boldsymbol{\rho}_{G}\left(V_{j}\right) \cup\left\{\left(V_{i}, V_{j}\right)\right\}, \boldsymbol{D}\right)-q\left(V_{j}, \boldsymbol{\rho}_{G}\left(V_{j}\right), \boldsymbol{D}\right)\)
        od;
        select the pair \(V_{i}, V_{j} \in \boldsymbol{V}\) for which \(\operatorname{diff}\left(V_{i}, V_{j}\right)\) is maximal;
        if \(\operatorname{diff}\left(V_{i}, V_{j}\right)>0\) then
            \(\boldsymbol{\rho}_{G}\left(V_{j}\right):=\boldsymbol{\rho}_{G}\left(V_{j}\right) \cup\left\{V_{i}\right\}\)
        fi
    until \(\operatorname{diff}\left(V_{i}, V_{j}\right) \leq 0\)
end
```

To conclude, we would like to point out that the learning algorithm as outlined above is not guaranteed to find a minimal I-map for the joint probability distribution underlying the process that generated the database used. It will be evident that one reason for the learning algorithm not finding such a digraph is that it incorporates a search heuristic that considers only a limited number of possible acyclic digraphs and therefore incidentally may skip the minimal I-maps. Yet another reason lies in the quality
measure used, however. For infinite size databases, the MDL-measure can be shown to prefer minimal I-maps over all other acyclic digraphs. Unfortunately, this property is not retained for finite size databases. So, by employing the MDL-measure for a finite size database, a digraph that is not a minimal I-map can be selected, even if every possible acyclic digraph were considered.

### 5.3 Assessing Probabilities

Only after the qualitative part of the Bayesian network in the making has been constructed and is considered robust, is its quantitative part specified. Specifying the quantitative part of a Bayesian network amounts to defining the probability assessment functions for the variables modelled in the network. The task of assessing all required probabilities tends to be by far the hardest task in Bayesian network building.

In order to facilitate the assessment task, approaches have been introduced that first attach qualitative signs to the arcs of a network's digraph, resulting in a so-called qualitative Bayesian network [Wellman, 1990]. The benefits are that these qualitative signs are easily elicited from domain expert, for example by using the method described by [Van der Gaag \& Helsper, 2004]. In addition, the signs pose constraints on the probability distributions to be assessed, which can be exploited upon further assessment; different combinations of qualitative signs, for example, can be an indication of disjunctive interaction patterns, such as the noisy-or described in Section 5.3.2 [Lucas, 2005]. In this section we will not focus on such qualitative approaches to quantification of Bayesian networks; the interested reader is referred to [Renooij, 2001].

### 5.3.1 Sources of Probabilistic Information

In most problem domains, at least some probabilistic information for Bayesian network quantification is readily available, be it from literature or from domain experts. The most common sources of probabilistic information are [Jensen, 1995]:

- literature on the domain of application (textbooks, journals, conferences);
- (statistical) data;
- (other) models of domain knowledge;
- interviews with domain experts.

Literature on the domain of application often provides abundant probabilistic information. Unfortunately, this information is very seldom directly amenable to encoding in the Bayesian network in the making: the information typically is not complete, it concerns variables that are not causally related, and so on. Medical literature, for example, often reports conditional probabilities of the presence of symptoms given a disease, but not always the probabilities of these symptoms occurring in the absence of disease; moreover, conditional probabilities for unobservable intermediate disease states are usually lacking. An additional, commonly found problem that prohibits direct use of probabilistic information from literature is that the characteristics of the population from which the information is derived, is not properly described or deviates seriously from the characteristics of the population for which the Bayesian network is being developed [Korver \& Lucas, 1993].

If the literature on the domain at hand does not provide for sufficient and reliable probability assessments for quantification of the network in the making, estimates may be obtained from statistical data or from other models of domain knowledge. Unfortunately, experience shows that even if comprehensive data collections and models are available, they very seldom contribute to the quantification task [Jensen, 1995, Korver \& Lucas, 1993]. In a medical data collection, for example, unobservable intermediate patho-physiological states are typically not recorded.

As a consequence, the majority of the probabilities required will have to be assessed by domain experts. The problems encountered when eliciting probabilities from experts are widely known [Tversky et al., 1982]; an expert's assessments may for example reflect various biases and may not be properly calibrated.

### 5.3.2 Simplifying Probability Assessment

The task of assessing all probabilities required for quantifying a Bayesian network can often be simplified by assuming the presence of one or more disjunctive interactions in the network [Pearl, 1988]. A disjunctive interaction pertains to two or more causes and their common effect: it specifies that any of these causes in itself suffices to cause the effect and that the likelihood of this cause causing the effect does not diminish when several other causes are present simultaneously.

More formally, we consider a set of statistical variables $\boldsymbol{V}$ and an acyclic digraph $G=\left(\boldsymbol{V}_{G}, \boldsymbol{A}_{G}\right)$ with $\boldsymbol{V}_{G}=\boldsymbol{V}$ where $\left(V_{1}, V_{0}\right), \ldots,\left(V_{m}, V_{0}\right) \in \boldsymbol{A}_{G}$, for some vertices $V_{0}, V_{1}, \ldots, V_{m}, m \geq 1$. Note that the digraph $G$ expresses that $V_{1}, \ldots, V_{m}$ model different causes of a common effect modelled by $V_{0}$. The variables $V_{1}, \ldots, V_{m}$ are said to exhibit a disjunctive interaction with respect to $V_{0}$ if the variables $V_{1}, \ldots, V_{m}$ adhere to the properties of accountability and exception independence. The property of accountability expresses that at least one of the causes modelled by $V_{1}, \ldots, V_{m}$ needs to be present for $V_{0}$ to adopt the value true, or alternatively, that the variable $V_{0}$ takes the value false whenever all its causes are absent. The property of accountability is stated in terms of probabilities as

$$
\operatorname{Pr}\left(v_{0} \mid \neg v_{1} \wedge \cdots \wedge \neg v_{m}\right)=0
$$

We now turn to the property of exception independence. This property is best explained in terms of a logical or-gate. Note that in a logical or-gate, the presence of any one of the causes $V_{1}, \ldots, V_{m}$ suffices for $V_{0}$ to adopt the value true. Now, if the logical relationship between $V_{0}$ and its separate causes is perturbed, we say that the logical or-gate is noisy. In a noisy or-gate, for each cause $V_{i}$ a so-called exception mechanism may inhibit the presence of the cause to result in $V_{0}=$ true; such an exception mechanism is often termed an inhibitor. The exception mechanisms in a noisy or-gate may be looked upon as statistical variables. If the exception mechanism for a cause $V_{i}$ has the value true, then it inhibits the occurrence of the effect $V_{0}=t r u e$; if $V_{i}$ 's exception mechanism has the value false, then it does not prevent the occurrence of the effect upon presence of the cause modelled by $V_{i}$. We therefore can model the exception mechanism of a cause $V_{i}$ by a statistical variable $I_{i}$ such that

$$
\operatorname{Pr}\left(v_{0} \mid \neg v_{1} \wedge \cdots \wedge \neg v_{i-1} \wedge\left(v_{i} \wedge i_{i}\right) \wedge \neg v_{i+1} \wedge \cdots \wedge \neg v_{m}\right)=0
$$

and

$$
\operatorname{Pr}\left(v_{0} \mid \neg v_{1} \wedge \cdots \wedge \neg v_{i-1} \wedge\left(v_{i} \wedge \neg i_{i}\right) \wedge \neg v_{i+1} \wedge \cdots \wedge \neg v_{m}\right)=1
$$

Exception independence for $V_{1}, \ldots, V_{m}$ now is the property that statistical variables $I_{1}, \ldots, I_{m}$ modelling the exception mechanisms for $V_{1}, \ldots, V_{m}$, respectively, are mutually independent. Because of the analogy outlined above, a disjunctive interaction is often coined a noisy or-gate.

The information that $V_{1}, \ldots, V_{m}$ exhibit a disjunctive interaction with respect to $V_{0}$ is not directly amenable to encoding in a Bayesian network. The information, however, imposes strong restrictions on the probability assessment function $\gamma_{V_{0}}$ for the vertex $V_{0}$ in the network. To support this observation, we consider the various values that have to be assessed to define the function $\gamma_{V_{0}}$, separately:

- for the value $\gamma_{V_{0}}\left(v_{0} \mid \neg v_{1} \wedge \cdots \wedge \neg v_{m}\right)$, we have

$$
\gamma_{V_{0}}\left(v_{0} \mid \neg v_{1} \wedge \cdots \wedge \neg v_{m}\right)=0
$$

from the property of accountability of a disjunctive interaction;

- for $i=1, \ldots, m$, let $I_{i}$ be the inhibitor for cause $V_{i}$ and let $\operatorname{Pr}\left(i_{i}\right)=q_{i}$; then, we have that

$$
\gamma_{V_{0}}\left(v_{0} \mid \neg v_{1} \wedge \cdots \wedge \neg v_{i-1} \wedge v_{i} \wedge \neg v_{i+1} \wedge \cdots \wedge \neg v_{m}\right)=1-q_{i}
$$

- let $c$ be an arbitrary configuration of the set of causes $\left\{V_{1}, \ldots, V_{m}\right\}$ and let $T_{c}=$ $\left\{i \mid c \wedge v_{i} \not \equiv\right.$ False $\} ;$ then, we have that

$$
\gamma_{V_{0}}\left(v_{0} \mid c\right)=1-\prod_{i \in T_{c}} q_{i}
$$

from the property of exception independence of a disjunctive interaction.
From these observations we have that, once causes $V_{1}, \ldots, V_{m}$ are known to exhibit a disjunctive interaction with respect to $V_{0}$, it suffices to assess only $n$ probabilities to define the probability assessment function $\gamma_{V_{0}}$ for vertex $V_{0}$.

Recall that a disjunctive interaction adheres to the properties of accountability and exception independence. We take a closer look at the property of accountability. Recall that this property states that all causes of an effect are known and explicitly modelled in the Bayesian network in the making. We observe that this property is hardly ever satisfied in real-life applications. Since incompleteness is inherent to any model of domain knowledge, a Bayesian network will never be complete: there will always be possible causes of an effect that are not explicitly modelled in the network. Unfortunately, if the property of accountability is not satisfied, it is not possible to exploit the model of the noisy or-gate as described above - not even it the property of exception independence is satisfied.

Consider the variables $V_{1}, \ldots, V_{m}$ modelling different causes of a common effect expressed by $V_{0}$. Now, assume that the property of exception independence is satisfied and that the property of accountability is not: we have that

$$
\operatorname{Pr}\left(v_{0} \mid \neg v_{1} \wedge \cdots \wedge \neg v_{m}\right)=p
$$

where $p>0$. To be able to nevertheless exploit the model of the noisy or-gate, we can extend the Bayesian network in the making with an extra statistical variable in such a way that the property of accountability is enforced and the property of exception
independence is retained. To this end, we introduce a new variable $V_{m+1}$ into the network as modelling an extra cause of the effect $V_{0}$ - this variable then is taken to capture all causes of $V_{0}$ that are not yet modelled in the original network. For this variable $V_{m+1}$, we then have

$$
\operatorname{Pr}\left(v_{0} \mid \neg v_{1} \wedge \cdots \wedge \neg v_{m} \wedge \neg v_{m+1}\right)=0
$$

and

$$
\operatorname{Pr}\left(v_{0} \mid \neg v_{1} \wedge \cdots \wedge \neg v_{m} \wedge v_{m+1}\right)=p
$$

Furthermore, we take $V_{m+1}$ to be independent of $V_{1}, \ldots, V_{m}$. Note that the introduction of the variable $V_{m+1}$ has not altered the meaning of the original causal mechanism. Also note that the property of accountability is now satisfied by $V_{1}, \ldots, V_{m+1}$ and that the model of the noisy or-gate applies.

Although the introduction of an extra variable as outlined before allows for exploiting the model of the noisy or-gate in the case where the property of accountability is not satisfied, adding an extra variable to the network is not satisfactory from a knowledge representational point of view: the additional variable has no well-defined meaning and therefore in a sense pollutes the Bayesian network. The model of the leaky noisy orgate now allows for capturing all yet unmodelled causes of a common effect implicitly [Henrion, 1989]. Consider once more the variables $V_{0}, V_{1}, \ldots, V_{m}$ and assume that the property of exception independence is satisfied and the property of accountability is not; recall that we have that

$$
\operatorname{Pr}\left(v_{0} \mid \neg v_{1} \wedge \cdots \wedge \neg v_{m}\right)=p
$$

where $p>0$. The probability $p$ is called a leak probability and is looked upon as expressing the probability that the effect $v_{0}$ occurs spontaneously. Now, by introducing an extra inhibitor $I_{0}$ with probability $\operatorname{Pr}\left(i_{0}\right)=q_{0}=1-p$ similar observations apply to the leaky noisy or-gate as for the noisy or-gate: the information that the variables $V_{0}, V_{1}, \ldots, V_{m}$ constitute a leak noisy or-gate again imposes strong restrictions on the definition of the probability assessment function $\gamma_{V_{0}}$ for the variable $V_{0}$ :

- for the probability $\gamma_{V_{0}}\left(v_{0} \mid \neg v_{1} \wedge \cdots \wedge \neg v_{m}\right)$, we have

$$
\gamma_{V_{0}}\left(v_{0} \mid \neg v_{1} \wedge \cdots \wedge \neg v_{m}\right)=p
$$

- for $i=1, \ldots, m$, let $I_{i}$ be the inhibitor for cause $V_{i}$ and let $\operatorname{Pr}\left(i_{i}\right)=q_{i}$; then, we have that

$$
\gamma_{V_{0}}\left(v_{0} \mid \neg v_{1} \wedge \cdots \wedge \neg v_{i-1} \wedge v_{i} \wedge \neg v_{i+1} \wedge \cdots \wedge \neg v_{m}\right)=1-q_{i}
$$

- let $c$ be an arbitrary configuration of the set of causes $\left\{V_{1}, \ldots, V_{m}\right\}$, let $T_{c}=\{i \mid$ $c \wedge v_{i} \not \equiv$ False $\}$, and let $q_{0}=1-p$; then, we have that

$$
\gamma_{V_{0}}\left(v_{0} \mid c\right)=1-q_{0} \cdot \prod_{i \in T_{c}}\left(\frac{q_{i}}{q_{0}}\right)
$$

from the property of exception independence of a disjunctive interaction.
Note that the values for the probability assessment function $\gamma_{V_{0}}$ are defined much in the same way as in the model of the noisy or-gate. Also note that, since in the Bayesian network in the making the effect of a cause cannot be isolated from the spontaneous occurrence of the effect, the probability assessments of the effect given a single cause are adjusted to account for this spontaneous leak before they are combined.

### 5.3.3 Eliciting Probabilities from Experts

The field of decision analysis offers various techniques for the elicitation of judgemental probabilities from experts [Von Winterfeldt \& Edwards, 1986, Morgan \& Henrion, 1990]. We briefly review the two techniques that are most often used for eliciting probabilities for Bayesian networks. The simplest technique is the use of a numerical probability scale. A probability scale is a horizontal or vertical line with the endpoints denoting a $0 \%$ and a $100 \%$ chance, respectively, and a few numerical anchors in between, for example to denote a $50 \%$ chance; Figure 5.2 illustrates the basic idea. For each probability required, a domain expert is asked to indicate his or her assessment on a separate scale. In communicating a probability to be assessed to a domain expert, the probability is often transcribed verbally in terms of frequencies. The expert is asked to envisage one hundred cases within a specific context and assess the number of cases that exhibit a certain characteristic. Experience shows that the use of a probability scale along with the frequency method provides experts little to go by and may result in highly inaccurate probability assessments [Van der Gaag et al., 1999].

A more elaborate technique for the elicitation of judgemental probabilities is the use of reference lotteries. A domain expert is presented with a choice between two lotteries, one of which pertains to the probability to be assessed and the other one serves as a reference. The reference lottery yields a desired reward with probability $p$ and a less desired outcome with probability $1-p$. The second lottery yields the same desired reward if a specific case exhibits a certain characteristic and the less desired outcome otherwise.

Example 5.3.1 Suppose that, in a medical domain of application, an expert has to assess the probability of a specific patient with metastatic cancer showing an increased level of serum calcium. The domain expert is presented, for instance, with a choice between a reference lottery and the lottery that yields $\$ 10,000$ if the patient upon examination shows an increased level of serum calcium and $\$ 1$ if the level of serum calcium is not increased in the patient. Figure 5.3 depicts this choice between lotteries.

The domain expert is asked to adjust the value of $p$ in the reference lottery until he or she is indifferent between the two lotteries. The resulting value of $p$ then is taken to be the conditional probability that had to be assessed. Experience shows that the use of reference lotteries for eliciting probabilities from domain experts may avert to at least some extent the problems of bias and poor calibration that are typically found in human probability assessment. The use of lotteries, however, tends to be quite time-consuming and, in fact, often turns out to be infeasible for probability elicitation for Bayesian networks as a result of the large size and complexity of a network in the making.

While the use of lotteries for probability elicitation on the one hand tends to be infeasible for quantifying a Bayesian network, the use of a probability scale on the other


Figure 5.2: A numerical scale for probability elicitation.


Figure 5.3: An example reference lottery.
hand tends to yield assessments that are too much inaccurate. The use of a probability scale, nonetheless, serves to yield an assessment, be it an inaccurate one, for every conditional probability required. These assessments may then be used as a starting point for further refinement. More details on probability elicitation methods, including a method designed by Utrecht University researchers, which is based on the use of a verbal-numerical probability scale, can be found in [Renooij, 2001b].

### 5.3.4 A Procedure for Probability Refinement

The assessments obtained for a Bayesian network are inevitably inaccurate, due to incompleteness of data and partial knowledge of the problem under study. Particularly assessments obtained from experts are known to be highly inaccurate [Tversky et al., 1982]. The inaccuracies in the probability assessments for a Bayesian network influence the reliability of the network's output. Sensitivity analysis is a general technique for studying the effects of the uncertainties in the parameters of a mathematical model on this model's outcome [Morgan \& Henrion, 1990]. For a Bayesian network, sensitivity analysis provides for example for studying the effects of the uncertainties in the assessments for the network's conditional probabilities on a probability of interest. There are various different types of sensitivity analysis. For a Bayesian network, the simplest type of sensitivity analysis amounts to systematically varying the assessment for one of the network's probabilities while keeping all other assessments fixed. Such an analysis serves to reveal the effect of just the conditional probability whose assessment is being varied, on a probability of interest. A sensitivity analysis in which a single assessment is varied, is termed a one-way sensitivity analysis. In a two-way sensitivity analysis of a Bayesian network, two probability assessments are varied simultaneously. In addition to the separate effects of variation of the two assessments, a two-way sensitivity analysis reveals the joint effect of their variation on a probability of interest. In essence, it is also possible to systematically vary more than two probability assessments at the same time. The results of such an analysis, however, are often hard to interpret.

Sensitivity analysis of a Bayesian network provides for studying the effects of the uncertainties in the various probability assessments of the network on a probability of interest. Studying these effects can to a large extent serve to support the elicitation of


Figure 5.4: A procedure for probability elicitation for Bayesian networks.
probabilities, as it gives detailed insight into the level of accuracy that is required for the various probabilities of the network, and as a result provides for focusing further elicitation efforts. We describe an elicitation procedure in which, alternatingly, sensitivity analyses are performed of a Bayesian network in the making and probability assessments are refined; this procedure is summarised in Figure 5.4. We elaborate on the various different steps of the procedure separately.

Step 1 In the first step of the elicitation procedure, initial assessments are acquired for all conditional probabilities for a Bayesian network in the making. As we have argued before, for most domains of application, experts will have to provide the majority of these initial assessments. In this phase of the construction of the Bayesian network, we apply an elicitation technique that aims at acquiring assessments within a short period of time. To this end, in a limited number of interview sessions, experts are asked to assess all probabilities required off the top of their heads, for example using a numerical, or verbal-numerical, probability scale. In addition, they are asked to indicate plausible intervals along with their assessments, that define a range of values in which the 'true' probability lies with reasonable certainty. These intervals should be pessimistic rather than optimistic to ensure that the uncertainties in the various different assessments are not underestimated. Since experts are thus allowed to express the uncertainties in their assessments, they will tend to be less reluctant to provide numerical statements than when they feel compelled to give exact numbers.

Instead of eliciting all probabilities required from domain experts, initial assess-
ments, for at least some of these probabilities, may be obtained from data, if available. If the data at hand is known to be imperfect, incomplete, or biased, then the assessments derived should be supplemented with fairly large plausible intervals to capture the uncertainties involved.

Step 2 The probability assessments obtained in the first step of the elicitation procedure will in general be highly uncertain and should have considerably large plausible intervals. For some of the probabilities, these initial assessments will nonetheless suffice. Other probabilities, however, will require assessments with a higher level of accuracy. The second step of our elicitation procedure is aimed at uncovering the latter probabilities. For this purpose, the Bayesian network in the making is subjected to various sensitivity analyses. In these analyses, every single probability assessment for the network is varied, as is every pair of assessments. These analyses require a lot of computation and generate a huge amount of data from which we need to decide which probability assessments require further attention. This computational and informational burden, including the technical details of sensitivity analysis, is further discussed in Section 6.1.

If performing all sensitivity analyses is too time-consuming to be practicable, the analyses can be restricted to a moderate number of assessments. To this end, experts may be asked to indicate the probability assessments that are the most likely to affect a probability of interest or, alternatively, to indicate a number of conceptually related probability assessments.

Step 3 In the second step of the elicitation procedure, various probability assessments have been identified that induce a considerably large effect on a probability of interest (see Section 6.1 on criteria for determining what 'large effect' means). These assessments are potential candidates for further refinement. In the third step of the procedure, the extent to which these assessments can actually be refined is determined; this depends on the current uncertainty in the assessment, indicated by the width of its plausible interval, and on the elicitation techniques used to arrive at the assessment. A probability assessment with a rather small plausible interval obtained from applying elaborate elicitation techniques may not lend itself to further refinement. An assessment that is not yet very certain, on the other hand, may be more easily improved upon. Refinement of such an uncertain assessment, however, is only actually possible if reliable sources of probabilistic information remain to be explored; examples of such sources of information include further literature search, the use of a panel of experts, the use of more elaborate elicitation techniques, and the study of data, for instance collected in a prospective study. From among the potential candidates for further refinement, therefore, the assessments are identified that have a considerable plausible interval and for which yet unexplored sources of probabilistic information are available.

Step 4 The probability assessments that have been selected in the third step of the elicitation procedure are assessments that can be refined. To actually refine these assessments, an investment of time and money is required. In the fourth step of the elicitation procedure it is investigated for each of the selected assessments whether refinement is cost-effective, by weighing the costs in terms of money and time to be invested with the benefits of higher accuracy. The benefits of having an assessment of higher accuracy in the Bayesian network in the making may be a higher accuracy of a computed probability of interest and an improved performance more in general. For
example, for a Bayesian network that is to be used for diagnostic purposes, performance may be measured as the percentage of correctly diagnosed cases. Refining a probability assessment for this network would only be worthwhile if it would increase the number of correct diagnoses (see Chapter 6 on evaluation). Once a Bayesian network in the making exhibits satisfactory overall behaviour, refining assessments may be found to be no longer cost-effective.

Recursive step The probability assessments that have been identified in the fourth step of the elicitation procedure are known to induce a considerable effect on a probability of interest; moreover, any such assessment can be cost-effectively refined. For these assessments, the entire elicitation procedure is recursively repeated. The probabilities concerned are assessed anew in the first step of the next cycle of the procedure, using yet unexplored sources of probabilistic information. Since the plausible intervals of the initial assessments for these probabilities do not underestimate the uncertainties involved, refinement is likely to result in smaller plausible intervals for the new assessments. In the second step, once again several sensitivity analyses are performed. These analyses should not only focus on the new assessments, but also on previously investigated assessments as their effect on a probability of interest upon variation may have changed as a result of the refinement of the network. In addition, analyses may be performed with respect to assessments that have not been investigated before. By recursively refining probability assessments, the performance of the Bayesian network in the making is likely to gradually improve. The elicitation procedure is stopped as soon as the costs of further elicitation outweigh the improvement in the network's performance or higher accuracy can no longer be attained due to lack of knowledge. You have now finished constructing the first prototype of your network and are ready for evaluating it practical use.

## Exercises

## Exercise 5.1

In constructing the qualitative part of a Bayesian network for a domain of application, a knowledge engineer aims at a digraph that is as sparse as possible. Give at least two reasons for this goal.

## Exercise 5.2

In the construction of the qualitative part of a Bayesian network for a domain of application, a cycle may arise. Describe at least two methods for removing the cycle from the network in the making.

## Exercise 5.3

Suppose that you are asked to construct a Bayesian network for a major medical center in Utrecht, to support specialists in the diagnosis and prognostication of various types of vascular disease. The network is to be constructed by hand, with the help of a domain expert.
a. The expert indicates that there are four important types of vascular disease; these types of disease are not mutually exclusive. Now, suppose that for pragmatical reasons, the network should have just one hypothesis variable. Explain how the various types of disease can be modelled in the network's digraph.
b. Suppose that you decide to model a single-valued discrete domain variable by means of two statistical variables. Explain how the relationship between these two variables can be modelled.
c. After the digraph of the network has been completed, it becomes clear that it requires simply too many probabilities for its quantification. Describe at least four ways to reduce the number of probabilities required. For each of these, describe in which situations it can lead to a substantial reduction of the number of probabilities required and in which situations it cannot.

## Exercise 5.4

Suppose that for the hospitals in the Netherlands, a Bayesian network is being built for the diagnosis of acute cardiac disorders and that the network is being handcrafted with the help of domain experts.
a. For eliciting the topology of the digraph of the network in the making, the concept of causality is used as a heuristic guiding rule. Give an example showing that the concept of causality is not always suitable for this purpose.

Suppose that the digraph of the Bayesian network in the making comprises the following subgraph:


For the variable Smoking the values true and false have been identified, denoted as s and $\neg s$, respectively; for the variable BloodPressure the values $b$ and $\neg b$ have been identified, for the variable Cholesterol the values $c$ and $\neg c$, and for the variable HeartAttack the values $h$ and $\neg h$.
b. Suppose that from an emergency medical center in New York, a data collection is available with the medical records of 12738 patients with heart disease. Explain why this data collection cannot be used just like that for the assessment of the probabilities required for the variable HeartAttack in the network fragment above.
c. Before commencing with the elicitation of the required probabilities, you would like to ask the domain expert whether or not a disjunctive interaction may be assumed for the variable HeartAttack and its parents. Which questions would you pose?

* d. Now suppose that the domain expert indicates that the variables Smoking, BloodPressure and Cholesterol satisfy the property of exception independence with respect to the variable HeartAttack. The domain expert further assesses the following probabilities:

$$
\begin{aligned}
& \operatorname{Pr}(h \mid \neg s \wedge \neg b \wedge \neg c)=0.05 \\
& \operatorname{Pr}(h \mid s \wedge \neg b \wedge \neg c)=0.6 \\
& \operatorname{Pr}(h \mid \neg s \wedge b \wedge \neg c)=0.8 \\
& \operatorname{Pr}(h \mid \neg s \wedge \neg b \wedge c)=0.9
\end{aligned}
$$

Is it possible to specify a complete probability assessment function for the variable HeartAttack on the basis of the available information? If no, then which information is missing; if yes, then give the complete assessment function.

* e. In practice, building a Bayesian network for a diagnostic application such as the one described above often turns out to be easier than the construction of a rulebased expert system for the same application. Give at least one explanation for this observation.


## * Exercise 5.5

Suppose that with the help of a domain expert, a digraph of a Bayesian network is constructed consisting of the variables $V_{1}, V_{2}, V_{3}$, and $V_{4}$, and the arcs $\left(V_{1}, V_{2}\right),\left(V_{2}, V_{3}\right)$, and $\left(V_{2}, V_{4}\right)$.

From the literature on the domain under study, the following probabilistic information is available about the four variables:

$$
\begin{aligned}
& \operatorname{Pr}\left(v_{1} \wedge v_{2}\right)=0.25 \\
& \operatorname{Pr}\left(\neg v_{1} \wedge \neg v_{2}\right)=0.3 \\
& \operatorname{Pr}\left(v_{2} \wedge \neg v_{3}\right)=0.2 \\
& \operatorname{Pr}\left(v_{2} \wedge v_{3}\right)=0.25 \\
& \operatorname{Pr}\left(\neg v_{2} \wedge v_{3}\right)=0.15 \\
& \operatorname{Pr}\left(v_{3} \mid \neg v_{1} \wedge v_{2} \wedge v_{4}\right)=0.4 \\
& \operatorname{Pr}\left(v_{4}\right)=0.8
\end{aligned}
$$

No further information had been found. Construct the assessment functions for as many variables as possible from the available information.

## Exercise 5.6

Consider the following procedure for eliciting probability assessments from domain experts:

- the expert is asked to use his own keywords and phrases to indicate his assessments for the various probabilities;
- the expert is then asked to rank order the keywords and phrases he used;
- the knowledge engineer subsequently associates numerical probabilities with the various keywords and phrases.

Describe the strengths and weaknesses of this procedure.

## Exercise 5.7

* a. Let $\boldsymbol{V}=\left\{V_{1}, V_{2}, V_{3}\right\}$ be a set of binary statistical variables. Let $\boldsymbol{D}$ be a database over $\boldsymbol{V}$ comprising the following cases:

$$
\begin{array}{ll}
v_{1} \wedge \neg v_{2} \wedge \neg v_{3} & \neg v_{1} \wedge v_{2} \wedge v_{3} \\
v_{1} \wedge v_{2} \wedge v_{3} & \neg v_{1} \wedge \neg v_{2} \wedge \neg v_{3} \\
v_{1} \wedge \neg v_{2} \wedge \neg v_{3} & v_{1} \wedge \neg v_{2} \wedge \neg v_{3} \\
v_{1} \wedge v_{2} \wedge v_{3} & \neg v_{1} \wedge \neg v_{2} \wedge v_{3} \\
v_{1} \wedge \neg v_{2} \wedge \neg v_{3} & \neg v_{1} \wedge \neg v_{2} \wedge \neg v_{3} \\
\neg v_{1} \wedge \neg v_{2} \wedge v_{3} & v_{1} \wedge v_{2} \wedge v_{3} \\
\neg v_{1} \wedge v_{2} \wedge \neg v_{3} & v_{1} \wedge v_{2} \wedge v_{3} \\
v_{1} \wedge \neg v_{2} \wedge \neg v_{3} & v_{1} \wedge v_{2} \wedge \neg v_{3} \\
v_{1} \wedge \neg v_{2} \wedge \neg v_{3} & v_{1} \wedge v_{2} \wedge v_{3} \\
v_{1} \wedge \neg v_{2} \wedge v_{3} & v_{1} \wedge v_{2} \wedge v_{3}
\end{array}
$$

There is no other information available on the variables $V_{1}, V_{2}$ and $V_{3}$.
Suppose that the database $\boldsymbol{D}$ is exploited for automated construction of a Bayesian network. For constructing the network, the B search heuristic is used in combination with the MDL quality measure. Now suppose that at some stage, the following digraph is constructed:


Compute the difference in quality that is achieved by adding the arc $\left(V_{1}, V_{2}\right)$ to this digraph.
b. In practical applications, the $B$ search heuristic is applied not only in combination with the MDL quality measure, but also in combination with other quality measures. An example of such a measure is the Akaike information criterion; this criterion is defined as

$$
Q_{A}(G, \boldsymbol{D})=\log P(G)-N \cdot H(G, \boldsymbol{D})-K
$$

where $G, \boldsymbol{D}, P, N, H$ and $K$ have the same meaning as in the $M D L$ quality measure. Now, for a given database $\boldsymbol{D}$, let $G_{A}$ be the digraph that is yielded by the $B$ search heuristic with the Akaike information criterion and let $G_{M D L}$ be the digraph that is yielded by the heuristic with the MDL quality measure. Describe in which respect the digraphs $G_{A}$ and $G_{M D L}$ will differ in general.

## Chapter 6

## Bringing Bayesian Networks into Practice

In various domains of application, ranging from medicine to meteorology, knowledgebased systems are being developed that build upon a Bayesian network for their knowledge representation. The previous chapter has dealt with the construction of such a Bayesian network. In this chapter, we discuss two issues that are concerned with enabling actual use of a Bayesian network in practice. Before a Bayesian network is introduced into practice, it should be analysed whether the network behaves as expected, and produces acceptable output. There are very few standard validity tests for Bayesian networks, especially for expert-elicited networks. Existing tests or analyses focus on the graphical structure (d-separation analyses), or on concepts such as sensitivity to findings and sensitivity to parameters [Jensen, Nielsen, 2007]. Recently, a validation framework was introduced that suggests to include more than above-mentioned criteria for reliability and predictive validity [Pitchforth \& Mengersen, 2013]. In Section 6.1, we will discuss parameter sensitivity analysis as a general method for studying a Bayesian network's robustness. In Section 6.2, we discuss the evaluation of the practical value of a Bayesian network, by providing some measures for establishing a network's quality.

As experience with applying the Bayesian network framework is building, it becomes apparent that, although the framework offers many advantages over earlier approaches to automated reasoning with uncertainty, it lacks with regard to intelligent control over reasoning. A Bayesian network provides for computing probabilities only, whereas in practice often decisions, such as what diagnostic tests to perform or what therapy to instil, are to be based on the computed probabilities. The provision of control over reasoning is generally considered one of the main contributions of artificial intelligence research to automated reasoning: knowledge-based systems thank their success to a large extent to their ability to apply specialised knowledge for pruning search spaces and for selectively gathering evidence. Since control over reasoning is a prerequisite for arriving at problem-solving behaviour that is satisfactory both from a computational and a user's point of view, it is not surprising that providing for control is an important issue in present-day Bayesian network research. In Section 6.3, we describe a problemsolving architecture that allows for automated control over reasoning.


$$
\begin{array}{ll}
\gamma(m c)= & 0.20 \\
& \\
\gamma(b \mid m c)= & 0.20 \\
\gamma(b \mid \neg m c)= & 0.05 \\
& \\
\gamma(c \mid b, i s c)= & 0.80 \\
\gamma(c \mid \neg b, i s c)= & 0.80 \\
\gamma(c \mid b, \neg i s c)= & 0.80 \\
\gamma(c \mid \neg b, \neg i s c)= & 0.05
\end{array}
$$

Figure 6.1: An example Bayesian network

### 6.1 Sensitivity Analysis

The reliability of the output of a Bayesian network can be investigated by studying its robustness. Robustness pertains to the extent to which the network's conditional probabilities influence the output when deviations from the specified assessments are assumed. In a medical application, for example, erroneous diagnoses or non-optimal treatment recommendations may result from building upon inaccurate assessments. For gaining detailed insight in output robustness, a Bayesian network can be subjected to a parameter sensitivity analysis. In the previous chapter, sensitivity analysis was mentioned as part of a probability elicitation procedure. Sensitivity analysis, in addition, can be used as a technique underlying parameter tuning, that is, changing probability assessments so that the network gives the desired output. Perhaps most importantly, sensitivity analysis can give insight into the range of probability assessments, or even evidence profiles, for which the outcome of the network is valid.

A sensitivity analysis of a mathematical model basically amounts to stepwise variation of one or more parameters and computing the output of the model in each step. This is a very demanding process from a computational point of view. Fortunately, the output probability of interest in a Bayesian network relates to the probability assessments, or parameters, varied by a simple mathematical function which can be computed quite efficiently from the network under study [Kjærulff \& Van der Gaag, 2000, Coupé \& Van der Gaag, 2002]. The form of this function heavily depends on how the variables corresponding to the parameters of interest and output of interest are located relative to each other in the network, as well as on the evidence entered. In this syllabus we restrict the discussion to one-way and two-way sensitivity analyse, where only one or two parameters, respectively, are varied simultaneously. We will consider the general form of the functions, followed by some properties, criteria for selecting interesting parameters, and examples.

As a running example we now consider the network in Figure 6.1, describing a piece of (fictitious and incomplete) medical information. The fragment describes the problems of metastatic cancer (denoted $M C$ ) for an arbitrary patient, with regard to the development of a brain tumour. Metastatic cancer may be detected by an increased level of serum calcium ( $I S C$ ). The presence of a brain tumour $(B)$ may be established from a CT scan $(C T)$. Severe headaches $(S H)$ are indicative of the presence of a brain tumour. Both a brain tumour and an increased level of serum calcium are likely to ultimately cause a patient to fall into a coma $(C)$. The strengths of these
dependences are described by the conditional probabilities. The probabilities specified for the variable $I S C$, for example, express that knowing whether or not metastatic cancer is present has a considerable influence on the probability of an increased level of serum calcium in an arbitrary patient. On the other hand, severe headaches are expressed as quite common in both patients with and without a brain tumour. Severe headaches thus have a low predictive value for a brain tumour. From the conditional probabilities specified for vertex $C$, we see that in the absence of both a brain tumour and an increased level of serum calcium, there is only a very small probability of a patient falling into a coma. The presence of either one of these causes in an arbitrary patient, however, suffices to render the probability of this patient falling into a coma in the near future quite high. Note that the two causes do not contribute to this probability independently: if one of the causes is present, then the presence of the other cause has no further influence on the probability of a patient falling into a coma. The two causes are said to exhibit a (negative) synergistic influence on their common effect.

### 6.1.1 What to Analyse?

A sensitivity analysis with respect to a prior probability of interest allows for assessing the quality and robustness of a Bayesian network in its reflecting a prior probability distribution for the domain of application. In the presence of case-specific observations, however, a Bayesian network may show very different sensitivities. To reveal these sensitivities, a sensitivity analysis may be performed with respect to a posterior, or conditional, probability. Such an analysis allows for validating the network's behaviour for specific cases or profiles, for example, profiles of typical patient populations in a medical application.

It is clear that a full-blown sensitivity analysis quickly becomes infeasible as it requires analysing the relation between each parameter, or combination of parameters, in the network, and every possible output probability of interest. The number of different output probabilities is exponential in the number of variables in the network. Fortunately, not every theoretically possible output probability will be of interest in the domain of application. In addition, for a given probability of interest, not every parameter has to be varied. In fact, only parameters pertaining to variables in the sensitivity set for a variable of interest need to be included in the sensitivity analysis. Given evidence $\boldsymbol{e}$ for a set of variables $\boldsymbol{E}$, the sensitivity set for a variable of interest $A$ contains all variables whose parameters may affect the posterior probability distribution for $A$ upon variation. More precisely, the sensitivity set is the set of variables $V \in \boldsymbol{V}$ for which none of the following holds:

- $V \notin \boldsymbol{\rho}_{G}^{*}(A)$ and $\boldsymbol{\sigma}_{G}^{*}(V) \cap \boldsymbol{E}=\emptyset$;
- $V \in \boldsymbol{\rho}_{G}^{*}(A)$ and $\left\langle\{V\} \cup \boldsymbol{\rho}_{G}(V)\right| \boldsymbol{E}|\{A\}\rangle_{G}^{d} ;$
- $V \notin \boldsymbol{\rho}_{G}^{*}(A),\left\langle\{V\} \cup \boldsymbol{\rho}_{G}(V)\right| \boldsymbol{E}|\{A\}\rangle_{G}^{d}$ and $\boldsymbol{\sigma}_{G}^{*}(V) \cap \boldsymbol{E} \neq \emptyset ;$

A sensitivity set clearly depends to a large extent on the case-specific observations that have been entered into the network. In our example Bayesian network, for instance, once the presence or absence of a metastatic cancer has been established in a patient, varying the probability assessments for the variable $I S C$ can no longer influence $\operatorname{Pr}(b)$.

The sensitivity set can be readily identified by inspecting the network's qualitative part; more details can be found in [Coupé \& Van der Gaag, 2002].

Upon varying a parameter, we have to make sure that the probabilities pertaining to the same conditional distribution from which we take the parameter continue to sum to one. Upon varying a single parameter, we adopt the standard assumption of proportional scaling: the other parameters are co-varied such that their mutual proportional relationship is kept constant.

### 6.1.2 One-way Sensitivity Analysis

For a Bayesian network, sensitivity analysis basically amounts to establishing, for each of the network's conditional probabilities, the sensitivity function that expresses a probability of interest in terms of the parameter under study.

In the sequel, we denote the probability of interest by $\operatorname{Pr}(A=a \mid \boldsymbol{e})$, or $\operatorname{Pr}(a \mid \boldsymbol{e})$ for short, where $a$ is a specific value of the variable $A$ of interest and $\boldsymbol{e}$ denotes the available evidence. The network's parameters are denoted by $x=\gamma\left(b_{i} \mid \boldsymbol{\rho}\right)$, where $b_{i}$ is a value of a variable $B$ and $\rho$ is a combination of values for the predecessors of $B$. We use $f_{\operatorname{Pr}(a \mid e)}(x)$ to denote the function that expresses the probability $\operatorname{Pr}(a \mid \boldsymbol{e})$ in terms of the parameter $x$; we often omit the subscript for the function symbol $f$, as long as ambiguity cannot occur.

## Defining the sensitivity function

Under the assumption of proportional scaling, any sensitivity function is a quotient of two functions that are (multi-)linear in the parameters under study. The numerator of the quotient describes the probability $\operatorname{Pr}(a \wedge \boldsymbol{e})$ as a function of the parameters and the denominator describes $\operatorname{Pr}(\boldsymbol{e})$ as a function of the parameters. More formally, in a one-way analysis, the function takes the form

$$
f(x)=\frac{c_{1} \cdot x+c_{2}}{c_{3} \cdot x+c_{4}}
$$

where the constants $c_{j}, j=1, \ldots, 4$, are built from the assessments for the numerical parameters that are not being varied.

If no evidence is entered into a Bayesian network, then the prior probability of interest relates linearly to any network parameter. Moreover, if the parameter under study pertains to a variable that is an ancestor of the variable of interest in the network's qualitative part, and the parameter's variable has no observed descendants, then the sensitivity function reduces to a linear function as well $\left(c_{3}=0\right)$. For any parameter associated with a variable outside the sensitivity set of the variable of interest, the sensitivity function is constant. We conclude that a sensitivity function is either a linear function or a fragment of a rectangular hyperbola:

$$
f(x)=\frac{r}{x-s}+t, \text { with } s=-\frac{c_{4}}{c_{3}}, \quad t=\frac{c_{1}}{c_{3}}, \quad \text { and } \quad r=\frac{c_{2}}{c_{3}}+s \cdot t
$$

A rectangular hyperbola has two branches and two asymptotes. Figure 6.2 illustrates the locations of the possible hyperbola branches relative to the two asymptotes. For $r<0$, the branches lie in the second (II) and fourth (IV) quadrants relative to the asymptotes $x=s$ and $f(x)=t$; for $r>0$, the branches are found in the first (I) and third (III) quadrants. Since any sensitivity function is continuous for $x \in[0,1]$,


Figure 6.2: Hyperbolas and their constants (the constraints on $s$ and $t$ are specific for sensitivity functions)
a hyperbolic sensitivity function is actually a fragment of one of the four possible hyperbola branches.

## Computing the sensitivity function

There basically exist two methods for computing the constants of a sensitivity function. In the first method, the constants are determined by computing from the network the probability of interest for up to three values for the parameter under study, and subsequently solving the resulting system of linear equations [Coupé \& Van der Gaag, 2002]; for the network computations, any of the standard propagation algorithms can be used. If a propagation algorithm is used that builds upon junction trees, then the constants of a sensitivity function can be obtained more efficiently [Kjærulff \& Van der Gaag, 2000].

Example 6.1.1 We illustrate performing a one-way sensitivity analysis of our example network. We begin by taking the prior $\operatorname{Pr}(c)$ for our probability of interest. We address the one-way analyses with respect to the parameters $\gamma(b \mid m c), \gamma(i s c \mid m c)$, and $\gamma(i s c \mid \neg m c)$, respectively. The results of these three prior analyses are shown in Figure 6.3. The figure displays, for example, the probability of interest $\operatorname{Pr}(c)$ as a


Figure 6.3: A one-way sensitivity analysis of the example Bayesian network.
function of the parameter $x=p(i s c \mid \neg m c)$, that is,

$$
\operatorname{Pr}(c)=0.57 \cdot x+0.21
$$

The general form of this formula follows, after marginalisation, from the factorisation of the joint distribution of our example network:

$$
\begin{aligned}
\operatorname{Pr}(c)= & \sum_{c_{M C}} \sum_{c_{B}} \sum_{c_{I S C}} \gamma\left(c \mid c_{B}, c_{I S C}\right) \cdot \gamma\left(c_{B} \mid c_{M C}\right) \cdot \gamma\left(c_{I S C} \mid c_{M C}\right) \cdot \gamma\left(c_{M C}\right) \\
= & \sum_{c_{M C}} \sum_{c_{B}} \gamma\left(c \mid c_{B}, i s c\right) \cdot \gamma\left(c_{B} \mid c_{M C}\right) \cdot \gamma\left(i s c \mid c_{M C}\right) \cdot \gamma\left(c_{M C}\right) \\
& +\sum_{c_{M C}} \sum_{c_{B}} \gamma\left(c \mid c_{B}, \neg i s c\right) \cdot \gamma\left(c_{B} \mid c_{M C}\right) \cdot \gamma\left(\neg i s c \mid c_{M C}\right) \cdot \gamma\left(c_{M C}\right) \\
= & \sum_{c_{B}} \gamma\left(c \mid c_{B}, i s c\right) \cdot \gamma\left(c_{B} \mid m c\right) \cdot \gamma(i s c \mid m c) \cdot \gamma(m c) \\
& +\sum_{c_{B}} \gamma\left(c \mid c_{B}, i s c\right) \cdot \gamma\left(c_{B} \mid \neg m c\right) \cdot \gamma(i s c \mid \neg m c) \cdot \gamma(\neg m c) \\
& +\sum_{c_{B}} \gamma\left(c \mid c_{B}, \neg i s c\right) \cdot \gamma\left(c_{B} \mid m c\right) \cdot \gamma(\neg i s c \mid m c) \cdot \gamma(m c) \\
& +\sum_{c_{B}} \gamma\left(c \mid c_{B}, \neg i s c\right) \cdot \gamma\left(c_{B} \mid \neg m c\right) \cdot \gamma(\neg i s c \mid \neg m c) \cdot \gamma(\neg m c) \\
= & \sum_{c_{B}} \gamma\left(c \mid c_{B}, i s c\right) \cdot \gamma\left(c_{B} \mid m c\right) \cdot \gamma(i s c \mid m c) \cdot \gamma(m c) \\
& +\sum_{c_{B}} \gamma\left(c \mid c_{B}, i s c\right) \cdot \gamma\left(c_{B} \mid \neg m c\right) \cdot x \cdot \gamma(\neg m c) \\
& +\sum_{c_{B}} \gamma\left(c \mid c_{B}, \neg i s c\right) \cdot \gamma\left(c_{B} \mid m c\right) \cdot \gamma(\neg i s c \mid m c) \cdot \gamma(m c) \\
& +\sum_{c_{B}} \gamma\left(c \mid c_{B}, \neg i s c\right) \cdot \gamma\left(c_{B} \mid \neg m c\right) \cdot(1-x) \cdot \gamma(\neg m c)
\end{aligned}
$$

You can now compute the constants directly from this expression and the given assessment functions. Alternatively, you can compute $\operatorname{Pr}(c)$ from the network for two different values of parameter $x$, e.g. $x_{1}$ and $x_{2}$ (note that this requires actually changing the assessment function $\gamma_{I S C}$ !), resulting in two different output probabilities $p_{1}$ and $p_{2}$; the constants $c_{1} \approx 0.57$ and $c_{2} \approx 0.21$ then follow by solving the following system of linear equations:

$$
\begin{aligned}
p_{1} & =c_{1} \cdot x_{1}+c_{2} \\
p_{2} & =c_{1} \cdot x_{2}+c_{2}
\end{aligned}
$$

Note that you choose the values of $x_{1}$ and $x_{2}$, and subsequently compute the probabilities $p_{1}$ and $p_{2}$ from the network. The above system therefore contains two expressions with two unknowns, $c_{1}$ and $c_{2}$, and can therefore be uniquely solved.

Next, we take for the probability of interest the posterior probability $\operatorname{Pr}(b \mid s h) . \mathrm{By}$ doing so, we assess the robustness of the diagnosis of a brain tumour for an arbitrary patient with a primary tumour who is suffering from severe headaches. We address the one-way analyses with respect to the parameters $\gamma(m c), \gamma(b \mid \neg m c)$, and $\gamma(s h \mid \neg b)$,


Figure 6.4: A one-way sensitivity analysis of the example Bayesian network.
respectively. The results of these posterior analyses are shown in Figure 6.4. The figure shows, for example, the probability of interest $\operatorname{Pr}(b \mid s h)$ as a function of the parameter $x=p(b \mid \neg m c)$, that is,

$$
\operatorname{Pr}(b \mid s h)=\frac{4 \cdot x+0.20}{x+3.80}
$$

Note that, in contrast with the prior analyses discussed before, the analyses for the posterior probability of interest reveal a non-linear relationship between the probability assessment that is being varied and the probability of interest.

## Selecting parameters deserving attention

Sensitivity analyses for a large number of parameters and several outputs of interest will typically result in a huge number of sensitivity functions that are to be examined. A number of criteria can be used to decide whether or not the effect of a parameter can be considerable enough to warrant further consideration, either for refining its assessment as discussed in Chapter 5, or for determining the implications of the analysis for the employability of the network in the domain of application. Here we will briefly discuss the following selection criteria: absolute effect, plausible effect, sensitivity value, vertex proximity, and admissible deviation.

> Absolute effect The absolute effect of a parameter on an output probability is simply the absolute difference $|f(0)-f(1)|$.

Example 6.1.2 We continue our example analysis. Figure 6.3(a) shows that varying the assessment for the probability $\gamma(b \mid m c)$ from 0 to 1 has a negligible effect on the probability of interest $\operatorname{Pr}(c)$ : the prior probability of a patient falling into a coma within the next three years increases from 0.31 to 0.34 , approximately. Figure 6.3(b) shows that varying the initial assessment for the probability $\gamma(i s c \mid m c)$ has a somewhat stronger effect on the probability of interest: $\operatorname{Pr}(c)$ now ranges from 0.22 to 0.34 . From Figure 6.3(c), to conclude, it is seen that varying the assessment for the probability $\gamma(i s c \mid \neg m c)$ has an even stronger effect on $\operatorname{Pr}(c)$ : the prior probability of a coma ranges from 0.21 to 0.78 . Note that the three analyses reveal the linear relationship between the probability assessment that is being varied and the probability of interest.

Plausible effect So far we have treated the probability assessments of our Bayesian network as exact point probabilities. As for most applications, however, the initially obtained assessments are quite uncertain. If this uncertainty is captured by supplementing each probability assessment with a plausible interval that defines a range of values in which the 'true' probability lies with reasonable certainty, then we can select parameters of interest based upon their plausible effect on the probability of interest. The plausible effect is now defined as the absolute effect within the plausible interval; the plausible effect is bounded by the absolute effect.

Example 6.1.3 In the prior analyses, for the three probability assessments under study in our example network, the plausible intervals are indicated in Figure 6.3 by shading. The figure shows that plausible variation of the parameter $\gamma(i s c \mid \neg m c)$ has the strongest effect on the probability of interest $\operatorname{Pr}(c)$. Varying the parameters $\gamma(b \mid m c)$ and $\gamma(i s c \mid m c)$, respectively, within their plausible intervals results in a rather small effect on the probability of interest. We recall from Figure 6.3 that the effect on $\operatorname{Pr}(c)$ of varying the assessment for $\gamma(b \mid m c)$ from 0 to 1 is smaller than the effect of varying $\gamma(i s c \mid m c)$ from 0 to 1 . By taking the plausible intervals into consideration, however, variation of the assessment for $\gamma(b \mid m c)$ has the stronger plausible effect. Especially since the plausible interval for this assessment is quite large, for example further elicitation efforts may better be directed at the probability $\gamma(b \mid m c)$ than at the probability $\gamma(i s c \mid m c)$.

Sensitivity value The sensitivity value of a parameter $x$ with respect to a probability of interest is defined as $\left|\frac{\partial f}{\partial x}\left(x_{0}\right)\right|$, the absolute value of the first derivative of the sensitivity function at the original value $x_{0}$ of the parameter. The sensitivity value thus captures the effect of infinitely small shifts in the parameter on the probability of interest.

Example 6.1.4 For the sensitivity function describing the posterior probability $\operatorname{Pr}(b \mid$ $s h$ ) as a function of the parameter $x=\gamma(b \mid \neg m c$ ) (see Figure 6.4(b)), we find for example that

$$
f^{\prime}(x)=\frac{0.384}{(0.16 \cdot x+0.608)^{2}}
$$

the parameter has an original value of 0.05 , so the sensitivity value for this parameter is $\left|f^{\prime}(0.05)\right|=1.01$. Alternatively, if we consider the effect of varying parameter $x=$ $\gamma(s h \mid \neg b)$ on our posterior probability of interest (Figure $6.4(\mathrm{c})$ ), then the sensitivity value equals

$$
\left.\left\lvert\, \frac{-0.059}{(0.92 \cdot 0.60+0.064)^{2}}\right.\right) \mid=0.155
$$

Vertex proximity The problem of using the sensitivity value as a measure of robustness, is that it often gives insight only in the effect of very small parameter variations. The effects of larger parameter shifts are of course captured by the absolute and plausible effects, but can also be studied by examining the vertex proximity.

The vertex of a hyperbola branch is the point $\left(x_{v}, f\left(x_{v}\right)\right)$ where $\left|f^{\prime}\left(x_{v}\right)=1\right|$. The proximity of a parameter's original value $x_{0}$ to the $x$-value of the hyperbola's vertex is an indication of possible sensitivity of the output of interest to variation of the parameter. The vertex-proximity can be easily computed from the constants of the sensitivity function:

$$
x_{v}= \begin{cases}s+\sqrt{|r|}, & \text { if } s<0 \\ s-\sqrt{|r|}, & \text { if } s>1\end{cases}
$$

Example 6.1.5 Again consider our example sensitivity functions in Figure 6.4. For the hyperbola branch describing the output probability as a function of parameter $\gamma(b \mid \neg m c)$ with original value 0.05 , we have that $x_{v}=-3.8+\sqrt{15}=0.07$; the vertex of the function is therefore very close to the parameter's original value, indicating that non-infinitesimal variation of the parameter could possibly have large effects on the output probability of interest. For parameter $\gamma(s h \mid \neg b)$, on the other hand, the $x$-value of the vertex is found around 0.19 , which can be considered quite distant from the parameter's original value of 0.60 . If the original value of the latter parameter had been 0.20 , however, then its sensitivity value ( 0.96 ) would perhaps deem the parameter irrelevant for further inspection, whereas the vertex-proximity would warn us for possibly significant effects of variation.

Admissible deviation An admissible deviation for a variable of interest and a given parameter, is a pair of real numbers $(\alpha, \beta)$ that describe the shifts to smaller values and to larger values, respectively, that are allowed in the parameter without inducing a change in the most likely value of the variable of interest. For a parameter with an original value of $x_{0}$, the admissible deviation $(\alpha, \beta)$ thus indicates that the parameter can be safely varied within the interval $\left[x_{0}-\alpha, x_{0}+\beta\right]$. To express that the parameter can be varied as far as the bounds of the probability interval, the symbol $\infty$ is used.

For establishing an admissible deviation, the intersections of a sensitivity function relating one value of the output variable of interest to a parameter, and those pertaining to the other values of the output variable, are computed. More specifically, the admissible deviation is established by computing the $x$-coordinates of the points where the sensitivity function $f_{\operatorname{Pr}\left(a_{i} \mid e\right)}(x)$ for some variable of interest $A$, intersects with the sensitivity functions $f_{\operatorname{Pr}\left(a_{j} \mid e\right)}(x), j \neq i$. We note that if the variable of interest has only two values, then the two sensitivity functions always intersect for $f(x)=0.5$. For socalled threshold decision making (see Section 6.3.1), instead of considering a change in most likely value, we can determine when some outcome probability of interest becomes smaller or larger than some pre-specified threshold; this basically amounts to establishing the intersections of a sensitivity function with the constant functions associated with the values of the threshold probabilities.

Example 6.1.6 We illustrate investigating the robustness of decisions by computing admissible deviations for our example Bayesian network. First, consider an arbitrary patient. The effects of varying parameter $x=\gamma(b \mid \neg m c)$ with original value $x_{0}=0.05$ on the probabilities of disease $\operatorname{Pr}(B)$ are shown in Figure 6.5(a) and are described by

$$
f_{\operatorname{Pr}(b)}(x)=0.80 \cdot x+0.04, \quad \text { and } \quad f_{\operatorname{Pr}(\neg b)}(x)=-0.80 \cdot x+0.96
$$

The sensitivity functions intersect for $x=0.575$, resulting in an admissible deviation of $(\infty, 0.525)$. The patient is most likely to not suffer from a brain tumour and this


Figure 6.5: The sensitivity functions for the two possible values of the variable of interest $B$ as a function of (a) parameter $\gamma(b \mid \neg m c)$, and (b) parameter $\gamma(s h \mid \neg b)$, given a severe headache.
diagnosis is quite robust to variation of the parameter under study. Now consider a patient with severe headaches. The effects of varying parameter $x=\gamma(s h \mid \neg b)$ with original value $x_{0}=0.60$ on the posterior probabilities of disease $\operatorname{Pr}(B \mid s h)$ are shown in Figure $6.5(\mathrm{~b})$ and are described by

$$
f_{\operatorname{Pr}(b \mid s h)}(x)=\frac{0.064}{0.92 \cdot x+0.064}, \quad \text { and } \quad f_{\operatorname{Pr}(\neg b \mid s h)}(x)=\frac{0.92 \cdot x}{0.92 \cdot x+0.064}
$$

The sensitivity functions intersect for $x=0.0696$, resulting in an admissible deviation of $(0.53, \infty)$. We conclude that the patient is most likely not to suffer from a brain tumour and this diagnosis again is quite robust to variation of the parameter under study.

### 6.1.3 Two-way Sensitivity Analysis

We now address a two-way sensitivity analysis of a Bayesian network. In a two-way sensitivity analysis, two probability assessments are varied simultaneously to reveal their joint effect on a probability of interest. Recall that under the assumption of proportional co-variation, any sensitivity function is a quotient of two functions that are (multi-)linear in the parameters under study. In a two-way analysis, the function is bi-linear and takes the general form

$$
f(x, y)=\frac{c_{1} \cdot x \cdot y+c_{2} \cdot x+c_{3} \cdot y+c_{4}}{c_{5} \cdot x \cdot y+c_{6} \cdot x+c_{7} \cdot y+c_{8}}
$$

where the constants $c_{j}, j=1, \ldots, 8$, are again built from the assessments for the numerical parameters that are not being varied.

As remarked before, two parameters can have a synergistic effect on the probability of interest; this means that the joint effect of varying the two parameters is different from the sum of their individual effects. Not every pair of varied parameters will have such an interaction effect on a probability of interest, however. For example, any two parameters that pertain to incompatible probabilities, in the sense of specifying complementary values for the same variable, will not interact. The function expressing the probability of interest in terms of two such assessments will lack a product term. A two-way sensitivity analysis involving assessments for incompatible probabilities does not reveal any unanticipated effects on a probability of interest beyond the effects shown
by one-way sensitivity analyses for the two assessments separately. Any such pair of assessments can therefore be excluded from the analysis.

Two-way sensitivity analysis in Bayesian networks has received far less attention by researchers than one-way analysis. No details are as yet known about the possible shapes of the two-way functions. It is obvious though that if either one of the parameters is fixed to an arbitrary value, then the two-way function degenerates to a one-way function in the other parameter, which is either monotonically increasing or monotonically decreasing; as a result, the minimum and maximum function values are still found for the extreme parameter values 0 and 1 . Research, in addition, has not addressed selection criteria. Here we will therefore limit our discussion to the selection criteria that trivially apply.

## Selecting parameters deserving attention

For selecting parameters from a two-way analysis that deserve further attention, we can use the absolute effect, the plausible effect, and the sensitivity value as before. In addition, we can use contour distance as a selection criterion.

Absolute and plausible effect The absolute effect of parameters $x$ and $y$ on a probability of interest is captured by the largest absolute difference found between the function values of the sensitivity function in the 'corners' $(0,0),(0,1),(1,0)$, and $(1,1)$ of the domain: $\max \{f(i, j)-f(k, l) \mid i, j, k, l \in\{0,1\}\}$. The plausible effect is defined analogously where the 'corners' of the domain are given by the plausible intervals for the two parameters.

Sensitivity value The sensitivity value of two parameters $x$ and $y$ with respect to a probability of interest can be defined in terms of a directional derivative (for details, see [Bolt \& Renooij, 2014]). Note that for a surface in 3D, the gradient in a certain point will typically depend on the direction under consideration, so we in fact have an infinite number of sensitivity values; the maximal gradient in a point, however, equals the length of the gradient vector in that point. The maximum 2-way sensitivity value at the original values $x_{0}$ and $y_{0}$ of the parameters therefore equals

$$
\left|\nabla f\left(x_{0}, y_{0}\right)\right|=\sqrt{\left(\frac{\partial f}{\partial x}\left(x_{0}, y_{0}\right)\right)^{2}+\left(\frac{\partial f}{\partial y}\left(x_{0}, y_{0}\right)\right)^{2}}
$$

Contour distance Two-way sensitivity functions are most easily interpreted using contour plots, in which contour lines connect the combinations of values for the two parameters that result in the same value for the probability of interest. The distance between two contour lines indicates the variation necessary in the two assessments to shift the probability of interest from one contour line to another. If the contour lines are very close to one another, then a small variation in the parameters under study suffices to have a strong effect on the probability of interest; if, in contrast, the contour lines are further apart, then the probability of interest is not very sensitive to variation of the two assessments.

Example 6.1.7 We illustrate performing a two-way sensitivity analysis for our example network. For our probability of interest, we once again take the prior probability


Figure 6.6: A two-way sensitivity analysis of the example Bayesian network.
$\operatorname{Pr}(c)$. We first address an analysis with respect to the parameters $x=\gamma(b \mid m c)$ and $y=\gamma(i s c \mid m c)$. The corresponding sensitivity function equals

$$
\operatorname{Pr}(c)=-0.15 \cdot x \cdot y+0.15 \cdot x+0.15 \cdot y+0.194
$$

From this function it is readily seen that the two probability assessments under study upon variation have a negative interaction effect on the probability of interest. The function is depicted in Figure 6.6. We observe from the figure that the distances between the contour lines differ, indicating that varying the parameters $\gamma(b \mid m c)$ and $\gamma(i s c \mid m c)$ simultaneously has a joint effect on the probability of interest $\operatorname{Pr}(c)$ beyond the effects of their separate variation; this joint effect is due to the synergistic influence of the variables $B$ and $I S C$ on the variable $C$ outlined before. We further observe that the contour lines are closer to one another in the lower left part of the figure than in the upper right part. If the assessments for the parameters $\gamma(b \mid m c)$ and $\gamma(i s c \mid m c)$ are both quite small, therefore, their variation will have a stronger effect on the probability of interest than if the initial assessments have a higher value. To variation within the plausible intervals of the assessments $\gamma(b \mid m c)=0.2$ and $\gamma(i s c \mid m c)=0.8$, as indicated by shading in Figure 6.6, the probability of interest shows a relatively low sensitivity. We further observe that the absolute effect of their joint variation on $\operatorname{Pr}(c)$ is not too strong.

Now we address an analysis pertaining to the assessments for the parameters $\gamma(b \mid$ $m c)$ and $\gamma(i s c \mid \neg m c)$; note that the two assessments under study pertain to incompatible probabilities, since they are conditioned on two different values of the same variable. The corresponding sensitivity function equals

$$
\operatorname{Pr}(c)=0.03 \cdot x+0.57 \cdot y+0.204
$$

This function is depicted in Figure 6.7. We observe from the figure that the contour lines, once again indicating values for the probability of interest $\operatorname{Pr}(c)$, are equidistant. Equidistance of contour lines indicates that simultaneously varying the probability


Figure 6.7: A two-way sensitivity analysis of the example Bayesian network.
assessments under study has no joint effect on the probability of interest beyond the effects of their separate variation. The two-way analysis therefore does not provide any information in addition to the information yielded by one-way analyses for the separate assessments.

To conclude, we performed a two-way sensitivity analysis of our example network with respect to the posterior probability of interest $\operatorname{Pr}(b \mid s h)$. We first address the analysis of varying the assessments for the parameters $x=\gamma(b \mid \neg m c)$ and $y=\gamma(s h \mid b)$ simultaneously. The corresponding two-way sensitivity function equals

$$
\operatorname{Pr}(b \mid s h)=\frac{1.10005 \cdot x \cdot y-0.00056 \cdot x+0.0559 \cdot y-0.00034}{x \cdot y-0.6268 \cdot x+0.0835 \cdot y+0.7811}
$$

In this function, the terms $-0.00056 \cdot x$ and $-0.6268 \cdot x$ pertain to the effect of variation of just the probability assessment $p(b \mid m c)$; the terms $0.0559 \cdot y$ and $0.0835 \cdot y$ pertain to the assessment $p(i s c \mid m c)$. The terms $1.10005 \cdot x \cdot y$ and $x \cdot y$ with each other capture the interaction effect of the two assessments on the network's probability of interest. These terms provide information that cannot be revealed by one-way analyses with respect to the two assessments separately. The function is depicted in Figure 6.8. Note that the contour lines are closest to one another in the lower right part of the figure, indicating a high sensitivity of the posterior probability of interest to high values for the probability $\gamma(b \mid \neg m c)$ and low values for the probability $\gamma(s h \mid b)$. For variation, within the plausible intervals of the initial assessments $\gamma(b \mid \neg m c)=0.05$ and $\gamma(s h \mid b)=0.8$, as indicated by shading in Figure 6.8, however, the probability of interest is relatively stable.

### 6.2 Evaluating Bayesian Networks

To establish its practical value, a real-life Bayesian network is typically subjected to an evaluation study using data from the domain of application. Such a study amounts


Figure 6.8: A two-say sensitivity analysis of the example Bayesian network.
to entering the data available for each problem case into the network and computing the most likely outcome. This outcome is then compared against a given standard of validity. The results of the study are often summarised in the percentage of correctly computed outcomes. This percentage correct is generally taken to convey the practical value of the network. For a medical diagnostic application, for example, a percentage correct of $85 \%$ is taken to indicate that the network is likely to establish the correct diagnosis for 85 out of every 100 patients. For many applications, this percentage would convey the information that the network performs quite satisfactorily.

Unfortunately, interpretation of the percentage correct of a Bayesian network is not as straightforward as is often suggested. The percentage should be interpreted with respect to a specific data collection. Now, each data collection is likely to include errors and to reflect the biases exhibited by the experts who collected the data. Moreover, the data will include the effects of random variation, especially in domains of a scientific nature. In the medical domain, for example, there is random variation in patient data, arising from biological differences between patients in the progression of pathological processes and from differences in the physicians' interpretation of symptoms and signs [Fletcher et al., 1996]. When two outcomes are almost equally likely for a patient, chance determines, to at least some extent, which outcome is entered into the patient's medical record as the most likely one. Random variation may thus affect a network's percentage correct, but the extent to which it does so is not expressed by the percentage.

Bayesian networks in essence do not yield a deterministic outcome. Rather, they produce a posterior probability distribution for their outcome variable. This distribution reflects the network's doubt as to the most likely outcome. Since the percentage correct of a Bayesian network does not take the computed posterior distribution into consideration, it does not reveal the extent of uncertainty in the outcome. To incorporate the network's doubt in the assessment of its practical value, evaluation scores from the field of statistical forecasting can be used. The use and interpretation of such a score is illustrated by means of an evaluation study of a real-life Bayesian network for the staging of oesophageal cancer.

### 6.2.1 The Percentage Correct and its Shortcomings

With the help of two experts in gastrointestinal oncology from the Netherlands Cancer Institute, Antoni van Leeuwenhoekhuis, a Bayesian network for the staging of oesophageal cancer has been constructed. The network captures the state-of-the-art knowledge about oesophageal cancer in the oesophagus network; for details on the network and its construction, see [Van der Gaag et al., 2002]. For studying the ability of the oesophagus network to correctly predict the stage of a patient's cancer, the medical records of 156 patients diagnosed with oesophageal cancer are available. For each patient, between 6 and 21 of the 25 symptoms and test results modelled in the network are available. For each patient, also the stage of his or her cancer, as established by the attending physician, is recorded. This stage can be either I, IIA, IIB, III, IVA, or IVB, in the order of advanced disease. The tumour's stage is indicative of the effects and complications to be expected from the different available therapeutic alternatives.

In general, to establish the practical value of a Bayesian network, for each case the outcome with highest posterior probability is determined from the network and compared against a given standard of validity. The results of such an evaluation study are summarised in the percentage of cases for which the network yields the correct outcome as the most likely one. To establish the practical value of the oesophagus network, for each patient all diagnostic symptoms and test results available were entered and the most likely stage for the patient's cancer was computed from the Bayesian network. The computed stage was then compared with the stage recorded in the data. The results from this comparison are shown in the matrix of Figure 6.9. The numbers on the diagonal of the matrix are the numbers of patients per stage for whom the network yields the same stage as the one recorded in the data. Taking the stages from the medical records as a standard of validity, we find that the network establishes the correct stage for 133 of the 156 patients, that is, the network has a percentage correct of $85 \%$.

As Bayesian networks in general, the oesophagus network in essence does not produce a deterministic outcome. Rather, it yields a probability distribution for its outcome variable. More specifically, it yields, for each patient, a posterior probability distribution over the six possible stages of his or her cancer. As an example, Figure 6.10 shows the distributions that are computed for three different real patients. Now, for some patients the computed posterior distribution clearly points to a single most likely

|  |  | I | IIA | IIB | network |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | III |  |  | IVA | IVB | total |
| data | I |  | 2 | 0 | 0 | 0 | 0 | 0 | 2 |
|  | IIA | 0 | 37 | 0 | 1 | 0 | 0 | 38 |
|  | IIB | 0 | 1 | 0 | 3 | 0 | 0 | 4 |
|  | III | 1 | 10 | 0 | 36 | 0 | 0 | 47 |
|  | IVA | 0 | 0 | 0 | 4 | 35 | 0 | 39 |
|  | IVB | 0 | 0 | 0 | 3 | 0 | 23 | 26 |
|  | total | 3 | 48 | 0 | 47 | 35 | 23 | 156 |

Figure 6.9: The results from the evaluation study, expressed in terms of the numbers of correctly and incorrectly staged patients.

| patient 1, stage IVA |  |
| :--- | :--- |
| stage I | 0 |
| stage IIA | 0 |
| stage IIB | 0.0159 |
| stage III | 0.0882 |
| stage IVA | 0.8245 |
| stage IVB | 0.0714 |

(a)

| patient 2, stage III |  |
| :--- | :--- |
| stage I | 0 |
| stage IIA | 0 |
| stage IIB | 0.0002 |
| stage III | 0.3616 |
| stage IVA | 0.3498 |
| stage IVB | 0.2884 |

(b)

| patient 3, stage III |  |
| :--- | ---: |
| stage I | 0.0222 |
| stage IIA | 0.3753 |
| stage IIB | 0.0459 |
| stage III | 0.3714 |
| stage IVA | 0.0916 |
| stage IVB | 0.0936 |

(c)

Figure 6.10: The posterior probabilities of the six stages for three different patients.
stage. For other patients, however, the posterior distribution can reveal considerable uncertainty. It is not unlikely, therefore, that incorrect conclusions of a network can be attributed to the effect of random variation, rather than to, for example, a modelling error. In the percentage correct reported for a network, however, the distribution of uncertainty over the various different outcomes is not taken into account. For example, for the patients shown in Figure 6.10, the oesophagus network's outcome is classified simply as correct for the first two patients and as incorrect for patient 3.

### 6.2.2 The evaluation score

As mentioned previously, Bayesian networks typically yield a probability distribution for their outcome variable. While for some cases from the domain of application the computed posterior distribution will point to a single most likely outcome, it may reveal considerable uncertainty for other cases. The percentage correct as a summary of evaluation results does not take these uncertainties into account. For assessing the practical value of a network, however, not just the most likely outcome but also the posterior distribution over the various possible outcomes should be taken into consideration.

Bayesian networks basically are probabilistic forecasters, as they repeatedly present predictions for an outcome variable in terms of probabilities. For the oesophagus network, for example, the posterior probability distribution that is computed for a specific patient can be looked upon as a forecast for the stage of this patient's cancer. Establishing the practical value of a Bayesian network thus amounts to assessing its quality as a forecaster. The quality of a probabilistic forecaster is often expressed in terms of its calibration, that is, the degree to which its forecasts match the true distribution of outcomes. For the oesophagus network, more specifically, we can say that it is (empirically) well calibrated if, among the patients for whom the network predicts a specific stage $S$ with probability $x_{S}$, the proportion of patients who in fact have stage $S$, denoted $x_{S}^{\prime}$, equals $x_{S}$. The smaller the difference between $x_{S}$ and $x_{S}^{\prime}$, that is, the closer the network's distribution matches the true distribution, the better calibrated the network is [Dawid, 1985, DeGroot \& Fienberg, 1983]. Building upon this concept of calibration, various scores for expressing the quality of a forecaster have been developed in the field of statistics.

Among the best-known evaluation scores is the Brier score [Panofsky \& Brier, 1968]. The basic idea of this score is illustrated for our oesophagus network. For each patient $i$, the network yields a forecast of posterior probabilities $p_{i j}$ over the stages $j=\mathrm{I}, \ldots$,

IVB. The Brier score $B_{i}$ of this forecast is defined as

$$
B_{i}=\sum_{j=\mathrm{I}, \ldots, \mathrm{IVB}}\left(p_{i j}-s_{i j}\right)^{2}
$$

where $s_{i j}=1$ if the medical record of patient $i$ states stage $j$, and $s_{i j}=0$ otherwise. If the network would yield the correct stage with certainty, that is, if the network would yield a correct deterministic forecast for the patient, then the associated Brier score would be equal to 0 . If the network would yield an incorrect deterministic forecast, the score would be 2 . For the forecast for a single patient, therefore, the Brier score ranges between 0 and 2, and the better the forecast, the lower the score. The Brier scores for the network's forecasts for the three patients from Figure 6.10 now are:

$$
B_{1}=0.04 \quad B_{2}=0.61 \quad B_{3}=0.56
$$

These scores reveal that the quality of the forecast for the first patient is very good, as expected. The other scores show that the forecasts for the patients 2 and 3 are of less quality. Recall that for patient 3 the forecast is unequivocal as a result of two stages being almost equally likely, among which is the correct stage. For patient 2 , there is even more uncertainty in the forecast, as there are three almost equally likely stages. These observations are reflected in the associated Brier scores: the score $B_{3}$ for patient 3 indicates higher quality than the score $B_{2}$ for the second patient. While in terms of the numbers of correctly and incorrectly staged patients the forecast for patient 2 is correct and the forecast for patient 3 is incorrect, the Brier score results in a more balanced and, hence, a more insightful quality assessment. Figure 6.11 summarises the Brier scores averaged over all, correctly and incorrectly staged, patients.

The quality of the oesophagus network as a forecaster can now be expressed in an overall score that is computed from the scores of the separate forecasts for our collection of patients. For $n$ patients, the overall Brier score $B$ is defined as

$$
B=\frac{1}{n} \sum_{i=1, \ldots, n} B_{i}
$$

It is readily seen that the overall Brier score again ranges between 0 and 2 , and the better the forecaster, the lower the score. For the oesophagus network, an overall Brier

|  |  |  |  |  | network |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- | :--- |
|  |  | I | IIA | IIB | III | IVA | IVB |
| data | I | $\mathbf{0 . 2 1}$ | - | - | - | - | - |
|  | IIA | - | $\mathbf{0 . 2 8}$ | - | 1.52 | - | - |
|  | IIB | - | 1.17 | - | 0.98 | - | - |
|  | III | 1.40 | 0.89 | - | $\mathbf{0 . 2 6}$ | - | - |
|  | IVA | - | - | - | 0.75 | $\mathbf{0 . 0 8}$ | - |
|  | IVB | - | - | - | 0.87 | - | $\mathbf{0 . 0 6}$ |

Figure 6.11: The results from the evaluation study, expressed in terms of average Brier scores.
score of 0.29 is found. To interpret this number, we compare the score with the overall scores obtained for two uninformed forecasters. The first forecaster gives a uniform probability distribution for each patient; this forecaster has an overall Brier score of 0.83. The second forecaster gives for each patient the prior distribution over the stages recorded in the data; this forecaster has an overall Brier score of 0.76 and is therefore slightly more informed than the uniform forecaster. The much lower Brier score of the oesophagus network now conveys the information that the network is quite informed and indeed builds upon its knowledge of oesophageal cancer to arrive at relatively good forecasts.

The percentage correct and the Brier score provide two ways of evaluating the practical value of a Bayesian network. The percentage correct treats outcomes as deterministic, whereas the Brier score takes the actual uncertainties into account.

### 6.3 A Problem-Solving Architecture

In this section, we describe a problem-solving architecture that allows for automated control over reasoning. The main purpose of exerting control over reasoning is to shape efficient and intelligent problem-solving behaviour. Exerting control involves monitoring and reflecting upon the reasoning process as it develops and taking decisions as to how it should proceed. To this end, strategic knowledge about the domain at hand is employed. As this knowledge may be non-probabilistic in nature, control over Bayesian network reasoning generally cannot be implemented in the framework in itself. The Bayesian network framework can therefore be embedded in a general problemsolving architecture [Van der Gaag \& Wessels, 1994a].

The Bayesian network problem-solving architecture is composed of two layers. The first layer offers the Bayesian network formalism and a variety of associated algorithms. The algorithms in this layer are characterised by their operating on a Bayesian network directly: various algorithms for probabilistic inference are comprised in the layer as are algorithms for example for reading independences from the qualitative part of a network. This layer is called the probabilistic layer of the architecture. The second layer of the problem-solving architecture is designed to provide for control over reasoning - it is termed the control layer. This layer offers a variety of methods for control for different types of problem solving and provides formalisms for representing the additional knowledge used by these methods. The layer for example can offer methods for selectively gathering evidence for diagnostic applications as well as methods for intelligently pruning and focusing Bayesian network inference. These control methods are the basic building blocks for shaping complex, domain-dependent problem-solving behaviour. The two layers of the architecture are strictly separated and communicate in a highly restricted fashion. The control layer queries the probabilistic layer for information about the represented joint probability distribution and the evidence entered so far, and, based upon this information, takes strategic decisions as to how to proceed. The probabilistic layer computes and returns the information it is asked for by the control layer.

The problem-solving architecture explicitly separates probabilistic reasoning from control over reasoning. Several advantages arise from such an explicit separation. A Bayesian network can be developed and refined, without being hampered by any algorithmic issues. Moreover, the representation of the joint probability distribution on the domain at hand is not obscured by non-probabilistic knowledge. In addition, a

Bayesian network can be re-used in different contexts for different purposes; a similar observation holds for the methods of control comprised in the control layer of the architecture. The architecture in addition provides for modelling decision problems. Knowledge about viable decisions and the preferences over their consequences involved in a decision problem are represented in the control layer, along with methods for solving the problem; the probabilistic layer comprises a Bayesian network that is used as a background knowledge base for providing the probabilities required. The idea of a meta-level problem-solving architecture pervades many areas of artificial intelligence research.

### 6.3.1 Threshold Decision Making

As an example of the type of problems that can be modelled using the problem-solving architecture, the threshold model for decision making is considered.

In the medical domain, Bayesian networks are often used for diagnostic purposes. A diagnostic Bayesian network typically comprises one or more variables modeling the presence of absence of disease, various variables modelling findings and results from diagnostic tests, and a number of intermediate variables modelling unobservable patho-physiological states. In the example network from Section 6.1 (Figure 6.1), for instance, the variable $B$ models the disease of interest, being the presence or absence of a brain tumour; the variable $M C$ models an unobservable state and the remaining variables capture findings and test results. A medical diagnostic Bayesian network is used for computing a most likely diagnosis for a patient given his or her presentation findings and test results.

The most likely diagnosis for a patient, along with its uncertainty, is generally taken by an attending physician to decide upon management of the patient. The physician may decide, for example, to start treatment right away. For the brain tumour example, the physician may decide to perform neurosurgery if a brain tumour is indicated. Alternatively, the physician may defer the decision whether or not to treat the patient until additional diagnostic information has become available, for example from a CT scan. Or, the physician may decide to withhold treatment altogether. To support choosing among these decision alternatives, the threshold model for patient management can be used. The threshold model for patient management is a typical example of something that can be implemented in the control layer of the problem-solving architecture.

The threshold model for patient management, or for decision making more in general, builds upon various threshold probabilities of disease [Pauker \& Kassirer, 1980].
The treatment threshold probability of disease, written $P^{*}(d)$ for disease $d$, is the probability at which an attending physician is indifferent between giving treatment and withholding treatment. If, for a specific patient, the probability of disease $\operatorname{Pr}(d)$ exceeds the treatment threshold probability, that is, $\operatorname{Pr}(d)>P^{*}(d)$, then the physician will decide to treat the patient as if the disease were known to be present with certainty. Alternatively, if $\operatorname{Pr}(d) \leq P^{*}(d)$, the physician will basically withhold treatment from the patient.

As a consequence of the uncertainty concerning the presence of disease in a patient, additional information from a diagnostic test may affect an attending physician's basic management decision. If the probability of disease exceeds the treatment threshold probability, then interpreting a negative test result may result in an updated probability of disease below the threshold probability. Alternatively, if the pretest probability of


Figure 6.12: The threshold model for patient management, indicating three threshold probabilities and the various decision alternatives at a physician's disposal.
disease falls below the treatment threshold probability, a positive test result may raise the probability of disease to a value above the threshold probability. To reckon with such effects, the threshold model for patient management includes another two threshold probabilities. The no treatment-test threshold probability of disease, written $P^{-}(d)$, is the probability at which the attending physician is indifferent between the decision to withhold treatment and the decision to obtain additional diagnostic information. The test-treatment threshold probability of disease, written $P^{+}(d)$, is the probability at which the physician is indifferent between obtaining additional information and starting treatment right away.

Figure 6.12 summarises the basic idea of the threshold model for patient management. As long as the diagnostic test under consideration has not been performed, a physician has three decision alternatives at his or her disposal. If the probability of disease $\operatorname{Pr}(d)$ for a patient falls below the no treatment-test threshold probability, that is, if $\operatorname{Pr}(d)<P^{-}(d)$, then the physician will withhold treatment from the patient without gathering additional diagnostic information. If the probability of disease exceeds the test-treatment threshold probability, that is, $\operatorname{Pr}(d)>P^{+}(d)$, then the physician will start treatment right away. Otherwise, that is, if $P^{-}(d) \leq \operatorname{Pr}(d) \leq P^{+}(d)$, the physician will perform the diagnostic test. After testing, there are only two decision alternatives left. If the updated probability of disease for the patient exceeds the treatment threshold probability, the physician will start treatment; otherwise treatment will be withheld from the patient.

The treatment threshold probability of disease $P^{*}(d)$ used in the threshold model is typically established by a physician after carefully weighing the various utilities involved. Utilities are numbers associated with a utility function and pertain to the presence or absence of disease on the one hand and giving or withholding treatment on the other hand. A utility function may be based on probabilistic information only and not involve any other information about the domain at hand. Yet, it may also incorporate non-probabilistic issues such as the cost of obtaining. From the expected utilities for giving and withholding treatment in view of the uncertainty concerning the presence of disease, the probability of disease at which the physician is indifferent between the two decision alternatives is readily determined; the basic idea is illustrated in Figure $6.13(\mathrm{a})$. For the brain tumour example, the physician will typically take into consideration the life expectancy for a patient, with and without a brain tumour, and the patient's attitude towards impaired health states; the physician can, for example, set the treatment threshold probability of a brain tumour at 0.15 . The two threshold probabilities $P^{-}(d)$ and $P^{+}(d)$ for deciding whether or not to perform a diagnostic test are established from the test's characteristics. For the brain tumour example, a possible diagnostic test is the CT scan. Suppose this test is added to the example network as a direct descendant of the variable $B$, together with the conditional probabilities

$$
\gamma_{C T}(c t \mid b)=0.95, \gamma_{C T}(c t \mid \neg b)=0.10
$$



Figure 6.13: The basic idea of establishing (a) the treatment threshold probability of disease, and (b) the no-treatment-test and test-treatment threshold probabilities.

The physician will typically weigh the discomfort of a CT scan for a patient against the additional information yielded by the scan; the physician, for example, may set the no treatment-test threshold probability of a brain tumour at 0.045 and the test-treatment threshold probability at 0.56 . The basic idea of establishing these two threshold probabilities is illustrated in Figure 6.13(b).

While the Bayesian network framework offers algorithms for computing the probability of disease and for processing evidence, it does not provide for computing the threshold probabilities and for comparing these to computed probabilities of disease: these tasks involve knowledge that cannot be expressed in the Bayesian network formalism and require computations beyond probabilistic inference. These tasks therefore are provided for by the control layer of the problem-solving architecture. The control layer in the problem-solving architecture can thus provide the additional information required to support the decision between treatment, no treatment, or deferring treatment until additional diagnostic information is available. If it is decided that additional diagnostic information is to be obtained, the control layer can also be used for determining what additional information is to be obtained. A simple method for selective gathering of such information is discussed next.

### 6.3.2 Selective Evidence Gathering

Another example of control over reasoning offered by the control layer of the problemsolving architecture, is a simple method for selective gathering of evidence for diagnostic problem solving with a Bayesian network.

In diagnostic problem solving, the objective is to identify a most likely explanation for a problem under consideration - this explanation then is the diagnosis of the problem. Establishing a diagnosis is achieved by gathering information about the manifestations of the problem at hand by applying tests to the problem. In most domains, it is not necessary to collect evidence on all possible manifestations before an accurate diagnosis is reached: information from only a few tests generally suffices. Moreover, it often is not desirable to apply all tests available as testing may be costly or damaging. In diagnostic problem solving, therefore, tests are not applied as a matter of course but instead are selected carefully. Selective evidence gathering, or test planning, now amounts to selecting the most useful tests to apply to a problem under consideration.

In essence, selective evidence gathering is concerned with three tasks. The first of
these is to select the test that is expected to yield the most useful information in the context of the evidence that is already available. When a test has been selected, the user, for example a physician, is requested to apply the test and to enter the evidence yielded. The second task of selective evidence gathering is to process this evidence. The third task is to decide whether enough evidence has been obtained as yet to confirm a diagnosis to sufficient extent. If still further information is required, the three tasks are executed recursively. We now take a closer look at these tasks in view of diagnostic problem solving with a Bayesian network. While the Bayesian network framework offers algorithms for computing probabilities and for processing evidence, thus providing for the second task of selective evidence gathering, it does not provide for valuing and selecting tests nor for deciding when to stop gathering information: these tasks involve knowledge that cannot be expressed in the Bayesian network formalism and require computations beyond probabilistic inference. These tasks therefore are provided for by the control layer of our problem-solving architecture.

In diagnostic problem solving, the variables discerned in the domain at hand play different roles; for example, some variables represent test outcomes, others represent unobservable, intermediate process states. For distinguishing between these different roles, the following types of vertex in the digraph of a Bayesian network are discerned [Henrion, 1989]: a hypothesis vertex represents one or more (mutually exclusive) hypotheses or disorders; an evidence vertex represents a variable whose value can be obtained by testing; all other vertices are intermediate vertices. The set of all vertices of the digraph $G$ of a Bayesian network can thus be partitioned into three mutually disjoint sets of vertices: $\boldsymbol{H}_{G}=\left\{H_{1}, \ldots, H_{n}\right\}, n \geq 1$, is the set of hypothesis vertices; the set of evidence vertices is denoted as $\boldsymbol{E}_{G}=\left\{E_{1}, \ldots, E_{m}\right\}, m \geq 1 ; \boldsymbol{I}_{G}$ is the set of intermediate vertices. The roles of the various vertices are modelled in the control layer of the problem-solving architecture and are not known to the probabilistic layer. In addition to knowledge concerning the roles of the vertices discerned, the control layer also specifies the additional knowledge required for assessing for each test the (expected) usefulness of information yielded by testing, and the extra knowledge involved in deciding when to stop gathering information.

For selective evidence gathering in diagnostic problem solving with a Bayesian network, generally two simplifying assumptions are made. First, a myopic approach to evidence gathering is taken, that is, evidence vertices to acquire information on are selected one by one. It is conceivable that in practical applications a non-myopic approach in which vertices are selected groupwise outperforms any method based on a myopic approach. Naively adopting a non-myopic approach, however, poses unsurmountable problems concerning computational complexity. The second simplifying assumption generally made is that the Bayesian network at hand comprises one hypothesis vertex $H$ only, that is, it is assumed that all hypotheses discerned in the domain are mutually exclusive. Note that this assumption prohibits reasoning about multiple interacting disorders. Relaxing this assumption and straightforwardly applying selective evidence gathering in view of a set of hypothesis vertices also causes serious computational problems, since then all possible combinations of values for all hypothesis vertices have to be considered. In the remainder of this section, we will equally take up the two assumptions mentioned above, although research results have indicated that the simplifying assumptions may be eased to some extent [Heckerman et al., 1993, Van der Gaag \& Wessels, 1994b, Sent, 2005].

Selective evidence gathering for diagnostic problem solving with a Bayesian network
may be envisioned as outlined below in pseudo-code. The evidence-gathering procedure takes the digraph $G$ of a Bayesian network and the set $\boldsymbol{E}$ of all yet uninstantiated evidence vertices for its input and yields as output a diagnosis $d$, which corresponds with a value assignment to the hypothesis variable.

```
procedure evidence-gathering \((G, \boldsymbol{E}, d)\)
    enough := false;
    while \(\boldsymbol{E} \neq \varnothing\) and not enough do
        dependent-vertices ( \(H, \boldsymbol{E}, \boldsymbol{E}^{\prime}\) );
        if \(E^{\prime} \neq \varnothing\) then
            select-vertex \(\left(\boldsymbol{E}^{\prime}, E_{j}\right)\);
            enter-evidence \(\left(E_{j}\right)\);
            \(\boldsymbol{E}:=\boldsymbol{E} \backslash\left\{E_{j}\right\} ;\)
            enough := verify-enough()
        else enough := true
    od;
    return diagnosis(d)
end
```

In principle, for selecting from a Bayesian network an appropriate vertex to acquire information on, each yet uninstantiated evidence vertex has to be examined as to the (expected) usefulness of information yielded. To this end, several probabilities are computed from the Bayesian network. These probabilities may reveal that for some of the uninstantiated evidence vertices, entering information has no influence whatsoever on the probabilities of the values of the hypothesis vertex and therefore is utterly useless in view of establishing a diagnosis. This property holds for all vertices that are independent of the hypothesis vertex given the evidence obtained so far. Now recall that the Bayesian network formalism allows for identifying independences from the digraph of a network without having to resort to probabilistic computations. In the main evidencegathering procedure, this property is exploited to save on the number of probabilities that has to be computed from the network. The dependent-vertices procedure is called upon to determine from a set $\boldsymbol{E}$ of uninstantiated evidence vertices the subset $\boldsymbol{E}^{\prime}$ of those vertices that are not d-separated from the hypothesis vertex $H$ given the evidence entered so far. For selecting an appropriate evidence vertex to acquire information on now only the vertices comprised in this set $\boldsymbol{E}^{\prime}$ are examined. Here, we will not further elaborate on the dependent-vertices procedure; for the computational issues involved, the reader is referred to [Geiger et al., 1990]. We would like to note that although the dependent-vertices procedure is called from the control layer of the problem-solving architecture, it itself is comprised in the probabilistic layer.

Once the set $\boldsymbol{E}^{\prime}$ of relevant uninstantiated evidence vertices has been determined, the select-vertex procedure selects from this set the evidence vertex to best acquire information on. Recently, the interest for selective evidence gathering - or test selection as it is most often referred to nowadays - has revived, resulting in a number of new measures and corresponding algorithms to select those vertices for which to acquire evidence (see e.g. [Sent, 2005]). Here, we will focus only on one of the earlier approaches where a utility function is used for discriminating between the various evidence vertices. This utility function assigns to each value of every evidence vertex a numerical quantity expressing the desirability, or utility, of obtaining this value. The utility functions in use for selective evidence gathering differ considerably in the way they value
information [Ben-Bassat, 1978, Glasziou \& Hilden, 1989]. For selecting an appropriate evidence vertex, the select-vertex procedure may employ any such utility function. As an example, we consider here a very simple utility function tailored to binary variables [Van der Gaag \& Wessels, 1994a]: the linear-value utility function $u$ is defined by

$$
u\left(E_{i}\right)=\left|\operatorname{Pr}\left(h \mid \widetilde{c}_{\boldsymbol{E}}\right)-\operatorname{Pr}\left(h \mid \widetilde{c}_{\boldsymbol{E}} \wedge E_{i}\right)\right|
$$

for each evidence vertex $E_{i} \in \boldsymbol{E}^{\prime}$, where $H$ is the hypothesis vertex and $\widetilde{c}_{\boldsymbol{E}}$ denotes the partial configuration of all evidence obtained so far. Note that for an uninstantiated evidence vertex $E_{i}$, the difference between $\operatorname{Pr}\left(h \mid \widetilde{c}_{\boldsymbol{E}}\right)$ and $\operatorname{Pr}\left(h \mid \widetilde{c}_{\boldsymbol{E}} \wedge e_{i}\right)$ indicates the confidence gained in the hypothesis $h$ if the evidence $E_{i}=$ true is observed; an analogous observation holds for the evidence $E_{i}=$ false. The usefulness of acquiring information on an evidence vertex, however, does not depend on a single value as it is uncertain which test result will be yielded for the problem at hand. For examining an evidence vertex, therefore, the utilities of its separate values are weighted with the probabilities that these values will be found. The result models the expected utility of acquiring information on the vertex. When employing the linear-value utility function $u$, the expected utility $\hat{u}$ for an evidence vertex $E_{i} \in \boldsymbol{E}^{\prime}$ is computed from

$$
\hat{u}\left(E_{i}\right)=\sum_{c_{E_{i}}} \operatorname{Pr}\left(c_{E_{i}} \mid \widetilde{c}_{\boldsymbol{E}}\right) \cdot u\left(c_{E_{i}}\right)
$$

To select the evidence vertex to best acquire information on, the select-vertex procedure computes the expected utilities for all relevant uninstantiated evidence vertices and then selects the vertex with highest expected utility.

Once an appropriate evidence vertex has been selected, the enter-evidence procedure prompts the user for a value for this vertex. The value obtained from the user is entered and subsequently processed in the Bayesian network at hand by the basic algorithms for probabilistic inference offered by the probabilistic layer of the problemsolving architecture. Note that the while-loop of the evidence-gathering procedure yields a sequence of prompts to the user concerning various evidence vertices.

The last task of selective evidence gathering is to investigate whether enough evidence has been collected to justify a decision to stop further gathering of information. For this task, a stopping criterion is employed. Such a stopping criterion may be based on several different principles. The principle of sufficiency of confirmation is to stop evidence gathering as soon as a diagnosis has been confirmed to sufficient extent by the available information: the probability of a (tentative) diagnosis is compared with a pre-set threshold value, and if this probability has surpassed the threshold value and is expected not to drop considerably, evidence gathering is stopped. Note that such a stopping criterion is based on probabilistic information only. Another principle a stopping criterion may be based upon is the principle of sufficiency of information. This principle is to stop evidence gathering if the expected utilities of all remaining evidence vertices have dropped below a pre-set threshold value; pursuing evidence gathering then is expected not to further contribute to establishing a diagnosis. A stopping criterion based on this principle may involve both probabilistic and non-probabilistic information from the domain at hand. In the evidence-gathering procedure the two principles are combined. The principle of sufficiency of information is seen in the condition of the main while-loop of the procedure: evidence gathering is stopped if all remaining uninstantiated evidence vertices are independent of the hypothesis vertex given the
evidence obtained so far. The verify-enough function completes the stopping criterion by implementing a test on sufficiency of confirmation.

## Exercises

## * Exercise 6.1

Consider the digraph of a Bayesian network with eight vertices $E_{1}, E_{2}, E_{3}, E_{4}, E_{5}, H_{1}$, $H_{2}$, and $H_{3}$, and nine arcs $\left(E_{1}, E_{2}\right),\left(E_{1}, H_{1}\right),\left(E_{1}, H_{2}\right),\left(E_{3}, E_{5}\right),\left(H_{1}, E_{2}\right),\left(H_{1}, E_{3}\right)$, $\left(H_{2}, E_{3}\right),\left(H_{2}, E_{4}\right)$, and $\left(H_{3}, E_{4}\right)$.
a. Establish the sensitivity set for the variable of interest $H_{1}$; also try to explain in an informal, intuitive way why certain variables are included in the sensitivity set and why other variables are not;
b. same question for $H_{1}$ given $E_{1}$;
c. same for $H_{1}$ given $E_{1}$ and $E_{3}$;
d. same for $H_{2}$ given $E_{1}, E_{2}, E_{3}, E_{4}$ and $E_{5}$;

Now suppose that an observation is entered for the variable $E_{2}$; the variable of interest is $H_{2}$ and we perform a sensitivity analysis. Assume that all variables are binary, that if some parameter probability $\operatorname{Pr}\left(v_{i} \mid c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)$ is varied then its complement $\operatorname{Pr}\left(\neg v_{i} \mid\right.$ $\left.c_{\boldsymbol{\rho}_{G}\left(V_{i}\right)}\right)$ is excluded from the analysis, and that we're interested only in the output probability $\operatorname{Pr}\left(h_{2}\right)$.
e. How many calls to a standard propagation algorithm for probabilistic inference are required for a sensitivity analysis of the network if the analysis is conducted straightforwardly, using 10 steps.?
f. How many of such network propagations are required if the functional form of the sensitivity function is exploited?

## Exercise 6.2

Consider the small Brain-tumour network from Figure 6.1.

* a. Suppose that only the variable $B$ has been observed. Which output probabilities are not influenced by variation of the parameters for $B$ ?
* b. Establish the minimal set of variables that need be observed to render variable $C T$ given $C$ insensitive to variation of the parameter probabilities of ISC.
* c. To which parameter probabilities does the output probability $\operatorname{Pr}(c \mid$ sh) relate linearly? Which parameter probabilities may, upon variation, have a non-linear effect on this output probability?
d. Analytically express the probability $\operatorname{Pr}(b)$ as a function of the parameter $p(b \mid m c)$.
e. Analytically express the probability $\operatorname{Pr}(c \mid b \wedge m c)$ as a function of the parameter $p(i s c \mid m c)$.


## Chapter 7

## Conclusions

Probabilistic reasoning with the Bayesian network framework is an exciting research area. First and foremost, the Bayesian network framework may be looked upon as a mathematically sound computational framework for probabilistic inference. From this point of view, we have addressed the algorithms offered by the framework. We have argued that although these algorithms have an exponential worst-case time complexity, they tend to behave polynomially for most real-life Bayesian networks. However, as applications of the framework grow larger, the Bayesian networks involved increase in size accordingly. Networks comprising hundreds or even thousands of vertices are no exception. For Bayesian networks of this size, the basic algorithms for probabilistic inference inevitably slow down problem solving despite their polynomial behaviour. Modern Bayesian network research therefore aims at developing more efficient algorithms. Efficiency is sought after in many different ways: existing algorithms are further optimised, new algorithms are designed for example based on simulation techniques, yet other optimisations derive from knowledge-based pruning and focusing. In addition, approximation and any-time algorithms are investigated.

The Bayesian network framework may also be looked upon as a framework for building knowledge-based systems. In fact, experience with developing applications of the framework is progressing rapidly. From this point of view, we have addressed the issue of building a Bayesian network for a domain of application. In many respects, building a Bayesian network resembles engineering a knowledge-based system more in general. Available knowledge-engineering methodologies, however, will have to be supplemented with methodologies tailored to Bayesian network building. Such methodologies are now emerging. As more and more data becomes available, the possibilities for automated construction are further explored, which is a popular topic of current research. As experience with applying the Bayesian network framework is building, the need for methods for knowledge-based control over reasoning is beginning to manifest itself. We have briefly addressed this issue and outlined shaping diagnostic problem solving with a Bayesian network.

## Chapter 8

## Solutions, Answers and Hints

This chapter presents solutions, answers or hints for selected exercises.

## Exercise 2.1

d. We prove the conditioning property stated in Proposition 2.2.8.

Let $\boldsymbol{V}$ be a set of statistical variables and let Pr be a joint probability distribution on $\boldsymbol{V}$. We have to show that

$$
\operatorname{Pr}(\boldsymbol{X})=\sum_{c_{\boldsymbol{Y}}} \operatorname{Pr}\left(\boldsymbol{X} \mid c_{\boldsymbol{Y}}\right) \cdot \operatorname{Pr}\left(c_{\boldsymbol{Y}}\right)
$$

for all sets of variables $\boldsymbol{X}, \boldsymbol{Y} \subseteq \boldsymbol{V}$. From the definition of conditional probability, we have that

$$
\begin{aligned}
\sum_{c_{\boldsymbol{Y}}} \operatorname{Pr}\left(\boldsymbol{X} \mid c_{\boldsymbol{Y}}\right) \cdot \operatorname{Pr}\left(c_{\boldsymbol{Y}}\right) & =\sum_{c_{\boldsymbol{Y}}} \frac{\operatorname{Pr}\left(\boldsymbol{X} \wedge c_{\boldsymbol{Y}}\right)}{\operatorname{Pr}\left(c_{\boldsymbol{Y}}\right)} \cdot \operatorname{Pr}\left(c_{\boldsymbol{Y}}\right)= \\
& =\sum_{c_{\boldsymbol{Y}}} \operatorname{Pr}\left(\boldsymbol{X} \wedge c_{\boldsymbol{Y}}\right)
\end{aligned}
$$

By employing the marginalisation property stated in Proposition 2.2.7, we further find that

$$
\begin{aligned}
\sum_{c_{\boldsymbol{Y}}} \operatorname{Pr}\left(\boldsymbol{X} \mid c_{\boldsymbol{Y}}\right) \cdot \operatorname{Pr}\left(c_{\boldsymbol{Y}}\right) & =\sum_{c_{\boldsymbol{Y}}} \operatorname{Pr}\left(\boldsymbol{X} \wedge c_{\boldsymbol{Y}}\right)= \\
& =\operatorname{Pr}(\boldsymbol{X})
\end{aligned}
$$

which concludes our proof.

## Exercise 2.4

To prove the property stated in the exercise, we show that

$$
\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Y} \wedge \boldsymbol{Z})=\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z}) \Leftrightarrow \operatorname{Pr}(\boldsymbol{X} \wedge \boldsymbol{Y} \mid \boldsymbol{Z})=\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z}) \cdot \operatorname{Pr}(\boldsymbol{Y} \mid \boldsymbol{Z})
$$

From the definition of conditional probability, we have that

$$
\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Y} \wedge \boldsymbol{Z})=\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z}) \Leftrightarrow \frac{\operatorname{Pr}(\boldsymbol{X} \wedge \boldsymbol{Y} \wedge \boldsymbol{Z})}{\operatorname{Pr}(\boldsymbol{Y} \wedge \boldsymbol{Z})}=\frac{\operatorname{Pr}(\boldsymbol{X} \wedge \boldsymbol{Z})}{\operatorname{Pr}(\boldsymbol{Z})}
$$

By division of both the left hand side and the right hand side of the latter expression by $\operatorname{Pr}(\boldsymbol{Z})$ and subsequent rearrangement of terms, we find that

$$
\begin{aligned}
& \operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Y} \wedge \boldsymbol{Z})=\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z}) \Leftrightarrow \\
& \quad \Leftrightarrow \frac{\operatorname{Pr}(\boldsymbol{X} \wedge \boldsymbol{Y} \wedge \boldsymbol{Z})}{\operatorname{Pr}(\boldsymbol{Z})}=\frac{\operatorname{Pr}(\boldsymbol{X} \wedge \boldsymbol{Z})}{\operatorname{Pr}(\boldsymbol{Z})} \cdot \frac{\operatorname{Pr}(\boldsymbol{Y} \wedge \boldsymbol{Z})}{\operatorname{Pr}(\boldsymbol{Z})} \Leftrightarrow \\
& \quad \Leftrightarrow \operatorname{Pr}(\boldsymbol{X} \wedge \boldsymbol{Y} \mid \boldsymbol{Z})=\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z}) \cdot \operatorname{Pr}(\boldsymbol{Y} \mid \boldsymbol{Z})
\end{aligned}
$$

which concludes our proof.

## Exercise 3.1

For $I_{\operatorname{Pr}}$ you need to use definitions based on $\operatorname{Pr}$.
a. See course handouts (slides).
b. A sketch of the proof is on the course handouts (slides); NB the full proof is quite complex!
c. We show that the independence relation $I_{\mathrm{Pr}}$ satisfies the property

$$
I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \rightarrow I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{W}, \boldsymbol{Y})
$$

for all mutually disjoint sets of variables $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W} \subseteq \boldsymbol{V}$, that is, we prove the third property stated in Theorem 3.1.2.

We assume that $I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W})$. From this observation, we have

$$
\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z} \wedge \boldsymbol{Y} \wedge \boldsymbol{W})=\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z})
$$

by definition, that is, we have that

$$
\frac{\operatorname{Pr}(\boldsymbol{X} \wedge \boldsymbol{Y} \wedge \boldsymbol{Z} \wedge \boldsymbol{W})}{\operatorname{Pr}(\boldsymbol{Y} \wedge \boldsymbol{Z} \wedge \boldsymbol{W})}=\frac{\operatorname{Pr}(\boldsymbol{X} \wedge \boldsymbol{Z})}{\operatorname{Pr}(\boldsymbol{Z})}
$$

From our assumption $I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W})$, we further have $I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{W})$ by the second property stated in the exercise. By definition, we therefore have that

$$
\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z} \wedge \boldsymbol{W})=\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z})
$$

that is,

$$
\frac{\operatorname{Pr}(\boldsymbol{X} \wedge \boldsymbol{Z} \wedge \boldsymbol{W})}{\operatorname{Pr}(\boldsymbol{Z} \wedge \boldsymbol{W})}=\frac{\operatorname{Pr}(\boldsymbol{X} \wedge \boldsymbol{Z})}{\operatorname{Pr}(\boldsymbol{Z})}
$$

Now consider the conditional probability $\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z} \wedge \boldsymbol{W} \wedge \boldsymbol{Y})$. By definition, we find that

$$
\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z} \wedge \boldsymbol{W} \wedge \boldsymbol{Y})=\frac{\operatorname{Pr}(\boldsymbol{X} \wedge \boldsymbol{Y} \wedge \boldsymbol{Z} \wedge \boldsymbol{W})}{\operatorname{Pr}(\boldsymbol{Y} \wedge \boldsymbol{Z} \wedge \boldsymbol{W})}
$$

From the previous observations, we further find that

$$
\begin{aligned}
\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z} \wedge \boldsymbol{W} \wedge \boldsymbol{Y}) & =\frac{\operatorname{Pr}(\boldsymbol{X} \wedge \boldsymbol{Z})}{\operatorname{Pr}(\boldsymbol{Z})}= \\
& =\frac{\operatorname{Pr}(\boldsymbol{X} \wedge \boldsymbol{Z} \wedge \boldsymbol{W})}{\operatorname{Pr}(\boldsymbol{Z} \wedge \boldsymbol{W})}= \\
& =\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z} \wedge \boldsymbol{W})
\end{aligned}
$$

From $\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z} \wedge \boldsymbol{W} \wedge \boldsymbol{Y})=\operatorname{Pr}(\boldsymbol{X} \mid \boldsymbol{Z} \wedge \boldsymbol{W})$, we have by definition that $I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z} \cup$ $\boldsymbol{W}, \boldsymbol{Y})$. We conclude that $I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \rightarrow I_{\operatorname{Pr}}(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{W}, \boldsymbol{Y})$.

## Exercise 3.2

We begin our proof by observing that, since $I$ is a semi-graphoid independence relation, it obeys the first four axioms stated in Definition 3.1.3. Now, we assume that $I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W})$ and $I(\boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W})$. We have that

$$
I(\boldsymbol{X}, \boldsymbol{Z}, \boldsymbol{Y} \cup \boldsymbol{W}) \rightarrow I(\boldsymbol{X}, \boldsymbol{Z} \cup \boldsymbol{W}, \boldsymbol{Y}) \rightarrow I(\boldsymbol{Y}, \boldsymbol{Z} \cup \boldsymbol{W}, \boldsymbol{X})
$$

by the weak union and symmetry axioms; in conjunction with our assumption $I(\boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W})$, we find

$$
I(\boldsymbol{Y}, \boldsymbol{Z} \cup \boldsymbol{W}, \boldsymbol{X}) \wedge I(\boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W}) \rightarrow I(\boldsymbol{Y}, \boldsymbol{Z}, \boldsymbol{W} \cup \boldsymbol{X}) \rightarrow I(\boldsymbol{X} \cup \boldsymbol{W}, \boldsymbol{Z}, \boldsymbol{Y})
$$

by the contraction and symmetry axioms.

## Exercise 3.3

For $I$ you need to use the four properties on which it is defined (Definition 3.1.3).

## Exercise 3.5

This exercise requires you to draw $16 D$-maps and $4 I$-maps.

## Exercise 3.7

a. The property holds (why?).
b. The property doesn't hold (why?).
c. The property $\left\langle\left\{V_{2}\right\}\right|\left\{V_{1}\right\}\left|\left\{V_{3}\right\}\right\rangle_{G}^{d}$ holds in the digraph $G$ since all chains in $G$ from $V_{2}$ to $V_{3}$ are blocked by the set of vertices $\left\{V_{1}\right\}$. For example, the chain $V_{2}, V_{1}, V_{3}$ from $V_{2}$ to $V_{3}$ is blocked by $\left\{V_{1}\right\}$ since $V_{1} \in\left\{V_{1}\right\}$; the chain $V_{2}, V_{5}, V_{3}$ from $V_{2}$ to $V_{3}$ is blocked by $\left\{V_{1}\right\}$ since $\left\{V_{5}, V_{6}\right\} \cap\left\{V_{1}\right\}=\varnothing$.
d. The property holds (why?).
e. The property $\left\langle\left\{V_{2}\right\}\right|\left\{V_{3}, V_{4}\right\}\left|\left\{V_{6}\right\}\right\rangle_{G}^{d}$ does not hold in the digraph $G$ since not every chain in $G$ from $V_{2}$ to $V_{6}$ is blocked by the set of vertices $\left\{V_{3}, V_{4}\right\}$. For example, the chain $V_{2}, V_{5}, V_{6}$ from $V_{2}$ to $V_{6}$ is not blocked by $\left\{V_{3}, V_{4}\right\}$.
f. The property doesn't hold (why?).

## Exercise 3.8

This exercise requires you to draw some of the $15 D$-maps and $63 I$-maps ( 24 with 6 arrows, 28 with 5,10 with 4 and 1 with 3 arrows).

## Exercise 3.9

A similar proposition for $u n$ directed $D$-maps and $I$-maps is proven in Lemma 3.2.6, as well as in the course handouts (slides).

## Exercise 3.10

Hint: it is easiest to try and use a graph as small as possible, with only three vertices.

## Exercise 3.13

Not necessarily (give an example!)

## Exercise 3.14

Hint: it is easiest to try and use a graph as small as possible; use three vertices for a) and c), four vertices for b) and d).

## Exercise 4.2

Your computations should give $\operatorname{Pr}\left(v_{1} \wedge v_{2} \wedge v_{3}\right)=0.09375 ; \operatorname{Pr}\left(v_{2} \wedge v_{3}\right)=0.24375$; $\operatorname{Pr}\left(v_{1} \mid v_{2} \wedge v_{3}\right)=0.38462 ; \operatorname{Pr}\left(v_{1} \vee v_{2} \vee \neg v_{3} \vee v_{4}\right)=0.985$.

## Exercise 4.3

Be inspired by Lemma 4.2 .8 or use the definition of $\lambda_{V_{i}}\left(V_{i}\right)$.

## Exercise 4.4

NB Contrary to the examples in Chapter 4, it is sufficient to illustrate the working of Pearl's algorithm by computing only those parameters required to establish the requested probabilities, rather than performing each and every computation done by the algorithm. See for example the answer to Exercise 4.5a.
a. Your computations should give: $\operatorname{Pr}\left(v_{4}\right)=0.56, \operatorname{Pr}\left(\neg v_{4}\right)=0.44$.
b. Your computations should give: $\operatorname{Pr}\left(v_{4} \mid v_{6}\right)=0.88, \operatorname{Pr}\left(\neg v_{4} \mid v_{6}\right)=0.12$.
c. Entering the evidence $V_{2}=$ false cannot influence the probabilities of the values of the vertex $V_{6}$ since this vertex is already instantiated. The evidence, however, may influence the probabilities of the values of each of the other vertices: neither of $V_{1}, V_{3}$, $V_{4}$, and $V_{5}$ is d-separated from $V_{2}$ by the set $\left\{V_{6}\right\}$, the set of vertices for which evidence has been entered.

## Exercise 4.5

a. For computing the probabilities $\operatorname{Pr}\left(v_{5}\right)$ and $\operatorname{Pr}\left(\neg v_{5}\right)$ of its values vertex $V_{5}$ applies the data fusion lemma:

$$
\begin{aligned}
& \operatorname{Pr}\left(v_{5}\right)=\alpha \cdot \pi_{V_{5}}\left(v_{5}\right) \cdot \lambda_{V_{5}}\left(v_{5}\right) \\
& \operatorname{Pr}\left(\neg v_{5}\right)=\alpha \cdot \pi_{V_{5}}\left(\neg v_{5}\right) \cdot \lambda_{V_{5}}\left(\neg v_{5}\right)
\end{aligned}
$$

Since no evidence has been entered in the network as yet, we have from Lemma 4.2.8 for the compound diagnostic parameter $\lambda_{V_{5}}$ of $V_{5}$ that $\lambda_{V_{5}}\left(v_{5}\right)=1$ and $\lambda_{V_{5}}\left(\neg v_{5}\right)=1$.

For the value $\pi_{V_{5}}\left(v_{5}\right)$ of its compound causal parameter, vertex $V_{5}$ computes

$$
\pi_{V_{5}}\left(v_{5}\right)=\gamma_{V_{5}}\left(v_{5} \mid v_{3}\right) \cdot \pi_{V_{5}}^{V_{3}}\left(v_{3}\right)+\gamma_{V_{5}}\left(v_{5} \mid \neg v_{3}\right) \cdot \pi_{V_{5}}^{V_{3}}\left(\neg v_{3}\right)
$$

where

$$
\begin{aligned}
\pi_{V_{5}}^{V_{3}}\left(v_{3}\right)= & \gamma_{V_{3}}\left(v_{3} \mid v_{1} \wedge v_{2}\right) \cdot \gamma_{V_{1}}\left(v_{1}\right) \cdot \gamma_{V_{2}}\left(v_{2}\right)+ \\
& +\gamma_{V_{3}}\left(v_{3} \mid \neg v_{1} \wedge v_{2}\right) \cdot \gamma_{V_{1}}\left(\neg v_{1}\right) \cdot \gamma_{V_{2}}\left(v_{2}\right)+ \\
& +\gamma_{V_{3}}\left(v_{3} \mid v_{1} \wedge \neg v_{2}\right) \cdot \gamma_{V_{1}}\left(v_{1}\right) \cdot \gamma_{V_{2}}\left(\neg v_{2}\right)+ \\
& +\gamma_{V_{3}}\left(v_{3} \mid \neg v_{1} \wedge \neg v_{2}\right) \cdot \gamma_{V_{1}}\left(\neg v_{1}\right) \cdot \gamma_{V_{2}}\left(\neg v_{2}\right)= \\
= & 0.2 \cdot 0.8 \cdot 0.5+0.6 \cdot 0.2 \cdot 0.5+0.5 \cdot 0.8 \cdot 0.5+0.1 \cdot 0.2 \cdot 0.5= \\
= & 0.35
\end{aligned}
$$

and

$$
\begin{aligned}
\pi_{V_{5}}^{V_{3}}\left(\neg v_{3}\right)= & \gamma_{V_{3}}\left(\neg v_{3} \mid v_{1} \wedge v_{2}\right) \cdot \gamma_{V_{1}}\left(v_{1}\right) \cdot \gamma_{V_{2}}\left(v_{2}\right)+ \\
& +\gamma_{V_{3}}\left(\neg v_{3} \mid \neg v_{1} \wedge v_{2}\right) \cdot \gamma_{V_{1}}\left(\neg v_{1}\right) \cdot \gamma_{V_{2}}\left(v_{2}\right)+ \\
& +\gamma_{V_{3}}\left(\neg v_{3} \mid v_{1} \wedge \neg v_{2}\right) \cdot \gamma_{V_{1}}\left(v_{1}\right) \cdot \gamma_{V_{2}}\left(\neg v_{2}\right)+ \\
& +\gamma_{V_{3}}\left(\neg v_{3} \mid \neg v_{1} \wedge \neg v_{2}\right) \cdot \gamma_{V_{1}}\left(\neg v_{1}\right) \cdot \gamma_{V_{2}}\left(\neg v_{2}\right)= \\
= & 0.8 \cdot 0.8 \cdot 0.5+0.4 \cdot 0.2 \cdot 0.5+0.5 \cdot 0.8 \cdot 0.5+0.9 \cdot 0.2 \cdot 0.5= \\
= & 0.65
\end{aligned}
$$

So,

$$
\pi_{V_{5}}\left(v_{5}\right)=0.6 \cdot 0.35+0.4 \cdot 0.65=0.47
$$

Analogously, vertex $V_{5}$ computes

$$
\pi_{V_{5}}\left(\neg v_{5}\right)=0.53
$$

By substituting the values $\lambda_{V_{5}}\left(v_{5}\right), \lambda_{V_{5}}\left(\neg v_{5}\right), \pi_{V_{5}}\left(v_{5}\right)$, and $\pi_{V_{5}}\left(\neg v_{5}\right)$ in the data fusion lemma and subsequently eliminating the normalisation constant $\alpha$ we find

$$
\begin{array}{ll}
\operatorname{Pr}\left(v_{5}\right) & =0.47 \\
\operatorname{Pr}\left(\neg v_{5}\right) & =0.53
\end{array}
$$

b. Your computations should result in $\operatorname{Pr}\left(v_{5} \mid v_{3}\right)=0.6, \operatorname{Pr}\left(\neg v_{5} \mid v_{3}\right)=0.4$.
c. Entering the evidence $V_{3}=$ true cannot influence the probabilities of the values of the vertex $V_{4}$. By exploiting the d-separation criterion it is easily seen that the vertex $V_{4}$ is independent of the vertex $V_{3}$ given the empty set $\varnothing$, that is, the set of vertices for which evidence has been entered: the only chain in the digraph $G$ between the vertices $V_{3}$ and $V_{4}$ is blocked by $\varnothing$. Since no other vertex in $G$ is d-separated from vertex $V_{3}$ by the empty set, the evidence for $V_{3}$ may influence the probabilities of the values of all other vertices in the network.

## Exercise 4.6

NB See hints for Exercises 4.4 and 4.5a: it is sufficient to compute only those parameters required to establish the requested probabilities.
a. $\operatorname{Pr}\left(v_{1} \mid v_{7}\right)=0.6, \operatorname{Pr}\left(\neg v_{1} \mid v_{7}\right)=0.4$.
b. Yes, possibly (why?).
c. This probability of a conjunction cannot be directly computed using Pearl's algorithm (why?). Explain how you would compute it (consider Pearl's algorithm as a black box).

## Exercise 4.7

NB See hints for Exercises 4.4 and 4.5a: it is sufficient to compute only those parameters required to establish the requested probabilities.
a. $\operatorname{Pr}\left(v_{3} \mid v_{1} \wedge \neg v_{6}\right)=0.6, \operatorname{Pr}\left(\neg v_{3} \mid v_{1} \wedge \neg v_{6}\right)=0.4$.
b. This probability of a conjunction cannot be directly computed using Pearl's algorithm (why?). Explain how you would compute it (consider Pearl's algorithm as a black box).
c. See b.

## Exercise 4.8

This can be proven using induction.

## Exercise 4.9

a. You should find either $\left\{V_{2}\right\}$ or $\left\{V_{3}\right\}$.
b. $\operatorname{Pr}\left(v_{5}\right)=0.34, \operatorname{Pr}\left(\neg v_{5}\right)=0.66$. NB See hints for Exercises 4.4 and $4.5 a$ : it is sufficient to compute only those parameters required to establish the requested probabilities.

## Exercise 4.10

a. $\left\{V_{2}\right\}$.
b. $\operatorname{Pr}\left(v_{3}\right)=0.3, \operatorname{Pr}\left(\neg v_{3}\right)=0.7$. NB See hints for Exercises 4.4 and $4.5 a$ : it is sufficient to compute only those parameters required to establish the requested probabilities.
c. These probabilities cannot be directly computed using Pearl's (data fusion) algorithm (why?). Explain how you would compute them (consider Pearl's algorithm as a black box).

## Exercise 4.11

a. You should find $\left\{V_{1}, V_{2}\right\}$ or $\left\{V_{2}, V_{3}\right\}$.

## Exercise 4.12

a. A possible loop cutset is $\left\{V_{2}, V_{3}, V_{7}, V_{8}\right\}$; you won't find $\left\{V_{10}\right\}$ (why?).

## Exercise 5.4

d. The available information suffices for specifying a complete probability assessment function for the variable HeartAttack. Since the property of exception independence is satisfied and sufficient probabilities have been assessed, the model of the leaky noisy or-gate can be exploited. The complete probability assessment function is defined by

$$
\begin{array}{ll}
\gamma_{H}(h \mid s \wedge b \wedge c) & =1-0.95 \cdot \frac{0.4}{0.95} \cdot \frac{0.2}{0.95} \cdot \frac{0.1}{0.95}=0.991 \\
\gamma_{H}(h \mid \neg s \wedge b \wedge c) & =1-0.95 \cdot \frac{0.2}{0.95} \cdot \frac{0.1}{0.95}=0.980 \\
\gamma_{H}(h \mid s \wedge \neg b \wedge c) & =1-0.95 \cdot \frac{0.4}{0.95} \cdot \frac{0.1}{0.95}=0.958 \\
\gamma_{H}(h \mid s \wedge b \wedge \neg c) & =1-0.95 \cdot \frac{0.4}{0.95} \cdot \frac{0.2}{0.95}=0.916 \\
\gamma_{H}(h \mid \neg s \wedge \neg b \wedge c) & =0.9 \\
\gamma_{H}(h \mid \neg s \wedge b \wedge \neg c) & =0.8 \\
\gamma_{H}(h \mid s \wedge \neg b \wedge \neg c) & =0.6 \\
\gamma_{H}(h \mid \neg s \wedge \neg b \wedge \neg c) & =0.05
\end{array}
$$

and the complementary function values.
e. Rule-based expert systems are briefly discussed in the Bachelor AI course "Kunstmatige Intelligentie".

## Exercise 5.5

Full assessment functions for all variables, except for variable $V_{4}$ can be established. Note the incoherence in the information!

## Exercise 5.7

a. The result of your computation, using ${ }^{10} \mathrm{log}$, should be -0.55 (what will happen as a result?).

## Exercise 6.1

a. $\left\{E_{1}, H_{1}\right\}$
b. $\left\{H_{1}\right\}$
c. $\left\{E_{3}, H_{1}, H_{2}\right\}$
d. $\left\{E_{2}, E_{3}, E_{4}, H_{1}, H_{2}, H_{3}\right\}$
e. 200 propagations are required if we assume 10 -step variation per function (i.e. each function is evaluated in 10 different points).
f. 25 propagations are required if we assume 3 propagations only for each non-linear function.

## Exercise 6.2

a. $\operatorname{Pr}(c t \mid B), \operatorname{Pr}(s h \mid B)$, and their complements.
b. $\{B\}$ (in addition to the already observed variable $C$ )
c. Linear: parameters of $I S C$ and $C$; non-linear: parameters of $M C, B$ and $S H$. Variation of parameters of $C T$ has no effect.

## Bibliography

[Andreassen et al., 1987] S. Andreassen, M. Woldbye, B. Falck, S.K. Andersen. MUNIN - A causal probabilistic network for interpretation of electromyographic findings. Proceedings of the Tenth International Joint Conference on Artificial Intelligence, 1987, pp. 366-372.
[Andreassen et al., 1991] S. Andreassen, R. Hovorka, J. Benn, K.G. Olesen, E.R. Carson. A model-based approach to insulin adjustment. In: M. Stefanelli, A. Hasman, M. Fieschi, J. Talmon. Proceedings of the Third Conference on Artificial Intelligence in Medicine. Lecture Notes in Medical Informatics 44, Springer Verlag, Berlin, 1991, pp. 239-248.
[Baioletti et al., 2011] M. Baioletti, G. Busanello, B. Vantaggi. Acyclic directed graphs representing independence models. International Journal of Approximate Reasoning, 52(1), 2011, pp. 2-18.
[Bellazzi et al., 1991] R. Bellazzi, C. Berzuini, S. Quaglini, D.J. Spiegelhalter, M. Leaning. Cytotoxic chemotherapy monitoring using stochastic simulation on graphical models. In: M. Stefanelli, A. Hasman, M. Fieschi, J. Talmon. Proceedings of the Third Conference on Artificial Intelligence in Medicine. Lecture Notes in Medical Informatics 44, Springer Verlag, Berlin, 1991, pp. $227-238$.
[Ben-Bassat, 1978] M. Ben-Bassat. Myopic policies in sequential classification. IEEE Transactions on Computers, vol. C-27, 1978, pp. $170-174$.
[Blanco et al., 2005] R. Blanco, I. Inza, M. Merino, J. Quiroga, P. Larranñaga. Feature selection in Bayesian classifiers for the prognosis of survival of cirrhotic patients treated with TIPS. Journal of Biomedical Informatics, vol. 38, 2005, pp. 376 388.
[Bolt \& Renooij, 2014] J.H. Bolt, S. Renooij. Local sensitivity of Bayesian networks to multiple simultaneous parameter shifts. Proceedings of the Seventh European Workshop on Probabilistic Graphical Models. Lecture Notes in Artificial Intelligence, vol. 8754, Springer Verlag, 2014, pp. $65-80$.
[Boneh, Nicholson \& Sonenberg, 2006] T. Boneh, A. Nicholson, L. Sonenberg. Matilda: a visual tool for modelling with Bayesian networks. International Journal of Intelligent Systems, vol. 21, 2006, pp. 1127 - 1150.
[Boose \& Gaines, 1988] J. Boose, B. Gaines. Knowledge Acquisition Tools for Expert Systems, Academic Press, London, 1988.
[Bouckaert, 1995] R.R. Bouckaert. Bayesian Belief Networks: from Construction to Inference, Ph.D. thesis, Utrecht University, 1995.
[Bruza \& Van der Gaag, 1994] P.D. Bruza, L.C. van der Gaag. Index expression belief networks for information disclosure. International Journal of Expert Systems: Research and Applications, vol. 7, 1994, pp. 107 - 138.
[Buchanan \& Shortliffe, 1984] B.G. Buchanan, E.H. Shortliffe. Rule-based Expert Systems. The MYCIN Experiments of the Stanford Heuristic Programming Project. Addison-Wesley, Reading, Massachusetts, 1984.
[Cheeseman, 1988] P. Cheeseman. An inquiry into computer understanding. Computational Intelligence, vol. 4, 1988, pp. $58-66$.
[Chen \& Pollino, 2012] S. Chen, C. Pollino. Good practice in Bayesian network modelling. Environmental Modelling and Software, vol. 37, 2012, pp. 134 - 145.
[Cooper, 1990] G.F. Cooper. The computational complexity of probabilistic inference using Bayesian belief networks. Artificial Intelligence, vol. 42, 1990, pp. $393-$ 405.
[Coupé \& Van der Gaag, 1998] V.M.H. Coupé and L.C. van der Gaag. Practicable sensitivity analysis of Bayesian belief networks. Proceedings of the Joint Session of the 6th Prague Symposium of Asymptotic Statistics and the 13th Prague Conference on Information Theory, Statistical Decision Functions and Random Processes, Union of Czech Mathematicians and Physicists, 1998, pp. 81-86.
[Coupé \& Van der Gaag, 2002] V.M.H. Coupé and L.C. van der Gaag. Properties of sensitivity analysis of Bayesian belief networks. Annals of Mathematics and Artificial Intelligence, vol. 36, 2002, pp. 323-356.
[Cousins et al., 1993] S.B. Cousins, W. Chen, M.E. Frisse. A tutorial to stochastic simulation algorithms for inference in belief networks. Artificial Intelligence in Medicine, vol. 5, 1993, pp. $315-340$.
[Cox, 1979] R.T. Cox. Of inference and inquiry - an essay in inductive logic. The Maximum Entropy Formalism, MIT Press, Cambridge, Massachusetts, 1970.
[Dagum \& Luby, 1993] P. Dagum, M. Luby. Approximating probabilistic inference in Bayesian belief networks is NP-hard. Artificial Intelligence, vol. 60, 1993, pp. 141 - 153.
[Dawid, 1985] A.P. Dawid. Calibration-based empirical probability. Annals of Statistics, vol. 13, 1985, pp. $1251-1274$.
[De Dombal et al., 1972] F.T. de Dombal, D.J. Leaper, J.R. Staniland, A.P. McCann, J.C. Horrocks. Computer-aided diagnosis of acute abdominal pain. British Medical Journal, vol. 2, 1972, pp. 9 - 13.
[De Dombal et al., 1974] F.T. de Dombal, D.J. Leaper, J.C. Horrocks, J.R. Staniland, A.P. McCann. Human and computer-aided diagnosis of abdominal pain: further report with emphasis on the performance of clinicians. British Medical Journal, vol. 4,1974 , pp. $376-380$.
[DeGroot \& Fienberg, 1983] M.H. DeGroot and S.E. Fienberg. The comparison and evaluation of forecasters. The Statistician, vol. 32, 1983, pp. $12-22$.
[Dempster et al., 1977] A.P. Dempster, N.M. Laird, D.B. Rubin. Maximum Likelihood from Incomplete Data via the EM Algorithm. Journal of the Royal Statistical Society, Series B 39(1): 1-38, 1977.
[Druzdzel \& Van der Gaag, 1995] M.J. Druzdzel, L.C. van der Gaag. Elicitation of probabilities for belief networks: combining qualitative and quantitative information. Proceedings of the Eleventh Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann, 1995, pp. 141-148.
[Druzdzel \& Van der Gaag, 2000] M.J. Druzdzel, L.C. van der Gaag. Building probabilistic networks: 'Where do the numbers come from?' Guest editors' introduction. IEEE Transactions on Knowledge and Data Engineering, 12(4):481-486, 2000.
[Finetti, 1970] B. de Finetti. Theory of Probability, Wiley, New York, 1970.
[Fletcher et al., 1996] R.H. Fletcher, S.W. Fletcher, and E.H. Wagner. Clinical Epidemiology. The Essentials, 3rd ed, Williams \& Wilkins, Baltimore, 1996.
[Friedman et al., 1997] N. Friedman, D. Geiger, M. Goldszmidt, G. Provan, P. Langley, P. Smyth. Bayesian Network Classifiers, Machine Learning, 1997, pp. 131-163.
[Geiger \& Pearl, 1988] D. Geiger, J. Pearl. On the logic of causal models, Proceedings of the Fourth Workshop on Uncertainty in Artificial Intelligence, 1988, pp. 136 147.
[Geiger et al., 1990] D.E. Geiger, T. Verma, J. Pearl. $d$-separation: from theorems to algorithms. In: M. Henrion, R.D. Shachter, L.N. Kanal, J.F. Lemmer. Uncertainty in Artificial Intelligence 5, Elsevier Science, Amsterdam, 1990, pp. 139 148.
[Glasziou \& Hilden, 1989] P. Glasziou, J. Hilden. Test selection measures. Medical Decision Making, vol. 9, 1989, pp. 133-141.
[Gorry \& Barnett, 1968] G.A. Gorry, G.O. Barnett. Experience with a model of sequential diagnosis. Computers and Biomedical Research, vol. 1, 1968, pp. 490 507.
[Guida \& Tasso, 1989] G. Guida, C. Tasso. Topics in Expert System Design, NorthHolland, Amsterdam, 1989.
[Heckerman et al., 1992] D.E. Heckerman, E.J. Horvitz, B.N. Nathwani. Toward normative expert systems. Part 1: The Pathfinder project. Methods of Information in Medicine, vol. 31, 1992, pp. $90-105$.
[Heckerman et al., 1993] D.E. Heckerman, E.J. Horvitz, B. Middleton. An approximate nonmyopic computation for value of information. IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 15, 1993, pp. 292 - 298.
[Helsper \& Van der Gaag, 2002] E.M. Helsper, L.C. van der Gaag. Building Bayesian networks through ontologies. In: F. van Harmelen (ed.). Proceedings of the 15th European Conference on Artificial Intelligence. IOS Press, Amsterdam, 2002, pp. $680-684$.
[Henrion, 1989] M. Henrion. Some practical issues in constructing belief networks. In: L.N. Kanal, T.S. Levitt, J.F. Lemmer (eds). Uncertainty in Artificial Intelligence 3, Elsevier Science, North-Holland, 1989.
[Jackson, 1990] P. Jackson. Introduction to Expert Systems. Addison-Wesley, Wokingham, 1990.
[Jensen, 1995] A.L. Jensen. Quantification experience of a DSS for mildew management in winter wheat. In: M.J. Druzdzel, L.C. van der Gaag, M. Henrion, F.V. Jensen (eds). IJCAI-95 Workshop on Building Probabilistic Networks: Where Do the Numbers Come From ?, 1995, pp. 23-31.
[Jensen, 1996] F.V. Jensen. An Introduction to Bayesian Networks. UCL Press, London, 1996.
[Jensen, Nielsen, 2007] F.V. Jensen \& T.D. Nielsen. Bayesian Networks and Decision Graphs, 2nd edition. Springer Verlag, 2007.
[Jensen et al., 1990] F.V. Jensen, J. Nielsen, H.I. Christensen. Use of Causal Probabilistic Networks as High Level Models in Computer Vision. Technical Report R-90-39, University of Aalborg, Denmark, 1990.
[Karp et al., 1989] R. Karp, M. Luby, N. Madras. Monte-Carlo approximation algorithms for enumeration problems. Journal of Algorithms, vol. 10, 1989, pp. 429 448.
[Kjærulff \& Van der Gaag, 2000] U. Kjærulff and L.C. van der Gaag. Making sensitivity analysis computationally efficient. Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann, San Francisco, 2000, pp. $317-325$.
[Korver \& Lucas, 1993] M. Korver, P.J.F. Lucas. Converting a rule-based expert system into a belief network. Medical Informatics, vol. 18, 1993, pp. $219-241$.
[Laskey \& Mahoney, 2000] K.B. Laskey, S. Mahoney. Network engineering for agile belief network models. IEEE Transactions on Knowledge and Data Engineering, vol. 12(4), 2000, pp. $487-498$.
[Lauritzen \& Spiegelhalter, 1988] S.L. Lauritzen, D.J. Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems. Journal of the Royal Statistical Society, Series B, vol. 50, 1988, pp. 157 -224 .
[Lucas, 2005] P.J.F. Lucas. Bayesian network modelling through qualitative patterns. Artificial Intelligence, vol 163, 2005, pp. 233 - 263.
[Lucas \& Van der Gaag, 1991] P.J.F. Lucas, L.C. van der Gaag. Principles of Expert Systems. Addison-Wesley, Wokingham, 1991.
[Morgan \& Henrion, 1990] M.G. Morgan and M. Henrion. Uncertainty, a Guide to Dealing with Uncertainty in Quantitative Risk and Policy Analysis. Cambridge University Press, Cambridge, 1990.
[Neapolitan, 2003] R.E. Neapolitan. Learning Bayesian Networks, Prentice Hall, 2003.
[Panofsky \& Brier, 1968] H.A. Panofsky and G.W. Brier. Some Applications of Statistics to Meteorology. The Pennsylvania State University, University Park, Pennsylvania, 1968.
[Pauker \& Kassirer, 1980] S.G. Pauker and J.P. Kassirer. The threshold approach to clinical decision making. New England Journal of Medicine, vol. 302, 1980, pp. 1109-1117.
[Pearl, 1988] J. Pearl. Probabilistic Reasoning in Intelligent Systems. Networks of Plausible Inference. Morgan Kaufmann, Palo Alto, 1988.
[Pearl et al., 1990] J. Pearl, D. Geiger, and T. Verma. The logic of influence diagrams, in: R.M. Oliver and J.Q. Smith (eds). Influence Diagrams, Belief Nets and Decision Analysis, John Wiley \& Sons, 1990, pp. $67-86$.
[Pearl \& Paz, 1985] J. Pearl, A. Paz. GRAPHOIDS: a graph-based logic for reasoning about relevance relations. In: B. Du Boulay, D. Hogg, L. Steels (eds). Advances in Artificial Intelligence 2, 1985, North-Holland.
[Pearl \& Verma, 1987] J. Pearl, T.S. Verma. The logic of representing dependencies by directed acyclic graphs, Proceedings of the Sixth National Conference on Artificial Intelligence, 1987, pp. 374-379.
[Peek \& Ottenkamp, 1997] N.B. Peek, J. Ottenkamp. Developing a decision-theoretic network for a congenital heart disease, in: E. Keravnou, C. Garbay, R. Baud J. Wyatt (eds). Proceedings of the 6th Conference on Artificial Intelligence in Medicine Europe. Springer-Verlag, Berlin, 1997, pp. 157-168.
[Pitchforth \& Mengersen, 2013] J. Pitchforth, K. Mengersen. A proposed validation framework for expert elicited Bayesian networks. Expert Systems with Applications, vol. 40, 2013, pp. 162-167.
[Pourret, Naim \& Marcot, 2008] O. Pourret, P. Naim, B. Marcot (editors). Bayesian Networks. A Practical Guide to Applications.Wiley, England, 2008.
[Renooij, 2001] S. Renooij. Qualitative Approaches to Quantifying Probabilistic Networks. Ph.D. Thesis, Institute of Information and Computing Sciences, Utrecht University, The Netherlands.
[Renooij, 2001b] S. Renooij. Probability elicitation for belief networks: Issues to consider. Knowledge Engineering Review, vol. 16, 2001, pp. 255-269.
[Sent, 2005] D. Sent. Test-selection Strategies for Probabilistic Networks. Ph.D. Thesis, Department of Information and Computing Sciences, Utrecht University, The Netherlands.
[Shafer, 1976] G. Shafer. A Mathematical Theory of Evidence, Princeton University Press, Princeton, 1976.
[Shafer \& Pearl, 1990] G. Shafer, J. Pearl. Readings in Uncertain Reasoning. Morgan Kaufmann, Palo Alto, 1990.
[Shortliffe \& Buchanan, 1984] E.H. Shortliffe, B.G. Buchanan. A model of inexact reasoning in medicine. In: B.G. Buchanan, E.H. Shortliffe. Rule-based Expert Systems. The MYCIN Experiments of the Stanford Heuristic Programming Project. Addison-Wesley, Reading, Massachusetts, 1984, pp. 233-262.
[Shwe et al., 1991] M.A. Shwe, B. Middleton, D.E. Heckerman, M. Henrion, E.J. Horvitz, H.P. Lehmann, G.F. Cooper. Probabilistic diagnosis using a reformulation of the INTERNIST-1/QMR knowledge base I: the probabilistic model and inference algorithms. Methods of Information in Medicine, vol. 30, 1991, pp. 241 -255 .
[Sommerville, 1992] I. Sommerville. Software Engineering, Addison-Welsy, Wokingham, 1992.
[Stahlschmidt et al., 2013] S. Stahlschmidt, H. Tausendteufel, W.K. Härdle. Bayesian networks for sex-related homicides: structure learning and prediction. Journal of Applied Statistics 40(6), 2013, pp. 1155-1171.
[Studený, 1989] M. Studený. Multiinformation and the problem of characterization of conditional-independence relations. Problems Control Information Theory, vol. 18, 1989, pp. $3-16$.
[Studený, 1992] M. Studený. Conditional independence relations have no finite complete characterization, in: S. Kubik, J.A. Visek (eds). Information Theory, Statistical Decision Functions and Random Processes, Kluwer, Dordrecht, 1992, pp. 377 - 396 .
[Studený, 1998] M. Studený. Bayesian networks from the point of view of chain graphs, in: G. F. Cooper, S. Moral (eds). Uncertainty in Artificial Intelligence. Proceedings of the 14th Conference, Morgan Kaufmann, San Francisco 1998, pp. 496503.
[Suermondt \& Cooper, 1990] H.J. Suermondt, G.F. Cooper. Probabilistic inference in multiply connected belief networks using loop cutsets. International Journal of Approximate Reasoning, vol. 4, 1990, pp. 283-306.
[Suermondt \& Cooper, 1991a] H.J. Suermondt, G.F. Cooper. Initialization for the method of conditioning in Bayesian belief networks. Artificial Intelligence, vol. 50, 1991, pp. $83-94$.
[Suermondt \& Cooper, 1991b] H.J. Suermondt, G.F. Cooper. A combination of exact algorithms for inference on Bayesian belief networks. International Journal of Approximate Reasoning, vol. 5, 1991, pp. $521-542$.
[Taroni et al., 2006] F. Taroni, C. Aitken, P. Garbolino, A. Biedermann. Bayesian Networks and Probabilistic Inference in Forensic Science, Wiley \& Sons, Chichester, 2006.
[Tversky et al., 1982] D. Kahneman, P. Slovic, and A. Tversky. Judgment under Uncertainty: Heuristics and Biases, Cambridge University Press, Cambridge, 1982.
[Van der Gaag, 1994] L.C. van der Gaag. A pragmatic view of the certainty factor model. The International Journal of Expert Systems: Research and Applications, vol. 7, 1994, pp. 289 - 300.
[Van der Gaag \& Helsper, 2002] L.C. van der Gaag, E.M. Helsper. Experiences with modelling issues in building probabilistic networks. In: A. Go'mez-Pe'rez and V.R. Benjamins (eds.). Knowledge Engineering and Knowledge Management: Ontologies and the Semantic Web. Proceedings of EKAW 2002. LNAI vol. 2473, Springer-Verlag, 2002, pp. 21-26.
[Van der Gaag \& Helsper, 2004] L.C. van der Gaag, E.M. Helsper. Defining classes of influences for the acquisition of probability constraints for Bayesian networks. In: R. Lo'pez de Ma'ntaras and L. Saitta (eds.). Proceedings of the 16th European Conference on Artificial Intelligence (ECAI 2004). IOS Press, Amsterdam, 2004, pp. 1101-1102.
[Van der Gaag \& Meyer, 1996] L.C. van der Gaag, J.-J.Ch. Meyer. Characterising normal forms for informational independence. Proceedings of IPMU'96: Information Processing and Management of Uncertainty in Knowledge-Based Systems, 1996, pp. $973-978$.
[Van der Gaag \& Renooij, 2001] L.C. van der Gaag, S. Renooij. Analysing sensitivity data from probabilistic networks. Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann, 2001, pp. 530-537.
[Van der Gaag et al., 1999] L.C. van der Gaag, S. Renooij, C.L.M. Witteman, B. Aleman, B.G. Taal. How to elicit many probabilities. Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence, Morgan Kaufmann, 1999, pp. 647-654.
[Van der Gaag et al., 2002] L.C. van der Gaag, S. Renooij, C.L.M. Witteman, B.M.P. Aleman, and B.G. Taal. Probabilities for a probabilistic network: A case-study in oesophageal cancer. Artificial Intelligence in Medicine, vol. 25, 2002, pp. 123 - 148.
[Van der Gaag \& Wessels, 1994a] L.C. van der Gaag, M.L. Wessels. Selective evidence gathering for diagnostic belief networks. AISB Quarterly, vol. 86, 1994, pp. 23 34.
[Van der Gaag \& Wessels, 1994b] L.C. van der Gaag, M.L. Wessels. Multiple-disorder diagnosis with belief networks. Proceedings of the Fifth Workshop on Principles of Diagnosis - DX'94, 1994, pp. 343-351.
[Van der Gaag \& Meyer, 1998] L.C. van der Gaag, J.-J.Ch. Meyer. Informational independence: models and normal forms. International Journal of Intelligent Systems, 13, 1998, pp. 83 - 109.
[Waal \& Van der Gaag, 2005] P.R. de Waal, L.C. van der Gaag. Stable Independence in Perfect Maps. Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence, AUAI Press, 2005.
[Warner et al., 1961] H.R. Warner, A.F. Toronto, L.G. Veasy, R. Stephenson. A mathematical approach to medical diagnosis: application to congenital heart disease. Journal of the American Medical Association, vol. 177, 1961, pp. 177 - 183.
[Wellman, 1990] M.P. Wellman. Fundamental concepts of qualitative probabilistic networks. Artificial Intelligence, vol. 44, 1990, pp. 257-303.
[Von Winterfeldt \& Edwards, 1986] D. von Winterfeldt, W. Edwards. Decision Analysis and Behavioral Research. Cambridge University Press, New York, 1986.


[^0]:    ${ }^{1}$ The reflexive closure of set $\boldsymbol{A}$ under $r$ is $r^{0}(\boldsymbol{A})=\boldsymbol{A}$, the transitive closure is $r^{+}(\boldsymbol{A})=r(\boldsymbol{A}) \cup$ $r^{+}(r(\boldsymbol{A}))$, and both combined gives $r^{*}(\boldsymbol{A})=r^{0}(\boldsymbol{A}) \cup r^{+}(\boldsymbol{A})$.

[^1]:    ${ }^{2}$ Note the difference between these propositions and the afore mentioned outcomes/values!

[^2]:    ${ }^{3}$ Writing $c_{\boldsymbol{V}}, \boldsymbol{V}=\left\{V_{1}, \ldots V_{m}\right\}$, instead of $V_{1}=$ some value $\wedge V_{2}=$ some value $\wedge \ldots \wedge V_{m}=$ some value is very convenient, especially if you don't care about the actual values and variables.

[^3]:    ${ }^{1}$ Note that these parameters are not the same as the parameters $\gamma$ specified by the network's assessment functions!

[^4]:    ${ }^{2}$ Note that in this case $\alpha$ cannot be trivially computed, since $\lambda_{V_{i}}^{V_{j_{k}}}\left(v_{j_{k}}\right)+\lambda_{V_{i}}^{V_{j_{k}}}\left(\neg v_{j_{k}}\right)$ typically does not equal 1. For this reason, in applying Pearl's algorithm, normalisation is usually postponed to the final step of data-fusion.

