Most Probable Explanations in Bayesian Networks: complexity and tractability

Johan Kwisthout

Model Based Systems Development Group,
Faculty of Science, Radboud University Nijmegen,
P.O. Box 9010, 6500GL Nijmegen, The Netherlands
johank@science.ru.nl

March 16, 2010

Abstract

An overview is given of definitions and complexity results of a number of variants of the problem of probabilistic inference of the most probable explanation of a set of hypotheses given observed phenomena.

1 Introduction

Bayesian or probabilistic inference of the most probable explanation of a set of hypotheses given observed phenomena lies at the core of many problems in diverse fields. For example, in a decision support system that facilitates medical diagnostics (like the systems described in [1], [2], [3], or [4]) one wants to find the most likely diagnosis given clinical observations and test results. In a weather forecasting system as in [5] or [6] one aims to predict precipitation based on meteorological evidence. But the problem is often also key in the computational models of economic processes [7, 8, 9], sociology [10], and cognitive tasks as vision or goal inference [11, 12]. Although these tasks may superficially appear different, the underlying computational problem is the same: given a probabilistic network, describing a set of stochastic variables and the (in)dependencies between them, and observations (or evidence) of the values for some of these variables, what is the most probable joint value assignment to (a subset of) the other variables?

Since probabilistic (graphical) models have made their entrance in domains like cognitive science (see e.g. the editorial of the special issue on probabilistic models of cognition in the *TRENDS in Cognitive Sciences* journal [13]), this problem now becomes more and more interesting for other investigators than those traditionally involved in probabilistic reasoning. However, the problem comes in many variants (e.g., with either full or partial evidence) and has many names (e.g., MPE, MPA, and MAP which may or may not refer to the same

problem variant) that may obscure the novice reader in the field. Apart from the naming conventions, even the question *how an explanation should be defined* depends on the author (compare e.g. the approaches in [14], [15], [16], and [17]). Furthermore, some computational complexity results may be counter-intuitive at first sight.

For example, finding the best (i.e., most probable) explanation is NP-hard and thus intractable in general, but so is finding a *good enough* explanation for any reasonable formalization of 'good enough'. So the argument that is is sometimes found in the literature (e.g. in [13]) and that can be paraphrased as "Bayesian abduction is NP-hard, but we'll assume that the mind approximates these results, so we're fine" is fundamentally flawed. However, when constraints are imposed on the structure of the network or on the probability distribution, the problem may become tractable. In order words: the optimization criterion is not a *source of complexity* [18] of the problem, but the network structure *is*, in the sense that unconstrained structures lead to intractable models in general, while imposing constraints to the structure sometimes leads to tractable models.

With this paper we intend to provide the computational modeler, who describes phenomena in cognitive science, economics, sociology, or elsewhere, an overview of complexity and tractability results in this problem, in order to assist her in identifying sources of complexity. An example of such an approach can be found in [19]. Here the Bayesian Inverse Planning model for goal inference [11] was examined and the conditions under which the model becomes intractable, respectively remains tractable were identified, allowing the modelers to investigate the (psychological) plausibility of these conditions.

While good introductions to explanation problems in Bayesian networks exist (see, e.g., [20] for an overview of explanation methods and algorithms), these papers appear to be aimed at the user-focused knowledge engineer, rather than at the computational modeler, and thus pay less attention to complexity issues. In this paper, we aim to bridge that gap, and focus on *tractability issues* in explanation problems, i.e., we address the question under which circumstances problem variants are tractable or intractable. We present definitions and complexity results related to Bayesian inference of the most probable explanation, including some new or previously unpublished results. The paper starts with some needed preliminaries from probabilistic networks, graph theory, and computational complexity theory. In the following sections the computational complexity of a number of problem variants is discussed. The final section concludes the paper and summarizes the results.

2   Preliminaries

In this section, we give a concise overview of a number of concepts from probabilistic networks, graph theory, and complexity theory, in particular definitions of probabilistic networks and treewidth, some background on complexity classes defined by probabilistic Turing Machines and oracles, and fixed-parameter tractability. For a more thorough discussion of these concepts, the

reader is referred to textbooks like [15], [21], [22], [23], [24], [25], and [26].

2.1   Bayesian Networks

A probabilistic network $\mathcal{B}$ is a graphical structure that models a set of stochastic variables, the (in-) dependencies among these variables, and a joint probability distribution over these variables. $\mathcal{B}$ includes a directed acyclic graph $\mathbf{G} = (\mathbf{V}, \mathbf{A})$, modeling the variables and (in-) dependencies in the network, and a set of parameter probabilities Γ in the form of conditional probability tables (CPTs), capturing the strengths of the relationships between the variables. The network models a joint probability distribution $\Pr(\mathbf{V}) = \prod_{i=1}^{n} \Pr(v_i \,|\, \pi(V_i))$ over its variables, where $\pi(V)$ denotes the parents of V in $\mathbf{G}$. We will use upper case letters to denote individual nodes in the network, upper case bold letters to denote sets of nodes, lower case letters to denote value assignments to nodes, and lower case bold letters to denote joint value assignments to sets of nodes. We will use $\mathbf{E}$ to denote a set of evidence nodes, i.e., a set of nodes for which a particular joint value assignment is observed.
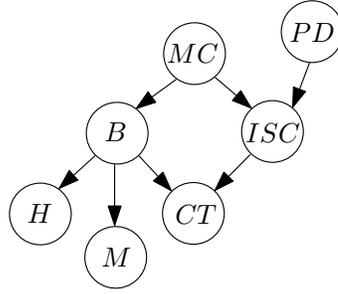
Throughout this chapter, we will refer to the *Brain tumor* network, shown in Figure 1, as a running example. This network, adapted from Cooper [27], captures some fictitious and incomplete medical knowledge related to metastatic cancer.. The presence of metastatic cancer (modeled by the node MC) typically induces the development of a brain tumor (B), and an increased level of serum calcium (ISC). The latter can also be caused by Paget's disease (PD). A brain tumor is likely to increase the severity of headaches (H). Long-term memory (M) is probably impaired, or even malfunctioning. Furthermore, it is likely that a CT-scan (CT) of the head will reveal a tumor if it is present, but it may also reveal other anomalies like a fracture or a lesion, which might explain an increased serum calcium.

Every (posterior) probability of interest in Bayesian networks can be computed using well known lemmas in probability theory, like Bayes' theorem ($\Pr(H \mid E) = \frac{\Pr(E|H)\Pr(H)}{\Pr(E)}$), marginalization ($\Pr(H) = \sum_{g_i} \Pr(H \wedge G = g_i)$), and the property that $\Pr(\mathbf{V}) = \prod_{i=1}^{n} \Pr(v_i \,|\, \pi(V_i))$. For example, from the definition of the Brain Tumor network we can compute that $\Pr(b \,|\, M = \mathrm{imp} \wedge CT = \mathrm{fract}) = 0.04$ and that $\Pr(mc \wedge \neg pd \,|\, M = \mathrm{norm} \wedge H = \mathrm{abs}) = 0.16$.

An important structural property of a probabilistic network is its *treewidth*. Treewidth is a graph-theoretical concept, which can be loosely described as a measure on the *locality* of the dependencies in the network. Formally, the treewidth of a probabilistic network, denoted by $\mathrm{tw}(\mathcal{B})$, is defined as the minimal width over all tree-decompositions of the moralization of $\mathbf{G}$. The moralization $\mathbf{M_G}$ of a directed graph $\mathbf{G}$ is the undirected graph, obtained by iteratively connecting the parents of all variables and then dropping the arc directions. The moral graph of the *Brain Tumor* network is shown in Figure 2.

A *tree-decomposition* of an undirected graph is defined as follows [22]:

Definition 2.1 (tree-decomposition). *A tree-decomposition of an undirected graph* $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ *is a pair* $\langle T, \mathcal{X} \rangle$*, where* $T = (I, F)$ *is a tree and* $\mathcal{X} = \{\mathbf{X_i} \,|\,$

$$
\begin{aligned}
\Pr(mc) &= 0.20 \\
\Pr(pd) &= 0.10 \\[4pt]
\Pr(b \mid mc) &= 0.20 \\
\Pr(b \mid \overline{mc}) &= 0.05 \\[4pt]
\Pr(M = \text{norm} \mid b) &= 0.50 \\
\Pr(M = \text{imp} \mid b) &= 0.40 \\
\Pr(M = \text{malf} \mid b) &= 0.10 \\
\Pr(M = \text{norm} \mid \overline{b}) &= 0.70 \\
\Pr(M = \text{imp} \mid \overline{b}) &= 0.25 \\
\Pr(M = \text{malf} \mid \overline{b}) &= 0.05 \\[4pt]
\Pr(H = \text{sev} \mid b) &= 0.70 \\
\Pr(H = \text{mod} \mid b) &= 0.25 \\
\Pr(H = \text{abs} \mid b) &= 0.05 \\
\Pr(H = \text{sev} \mid \overline{b}) &= 0.30 \\
\Pr(H = \text{mod} \mid \overline{b}) &= 0.20 \\
\Pr(H = \text{abs} \mid \overline{b}) &= 0.50
\end{aligned}
$$

$$
\begin{aligned}
\Pr(isc \mid mc, pd) &= 0.95 \\
\Pr(isc \mid mc, \overline{pd}) &= 0.80 \\
\Pr(isc \mid \overline{mc}, pd) &= 0.70 \\
\Pr(isc \mid \overline{mc}, \overline{pd}) &= 0.20 \\[4pt]
\Pr(CT = \text{tum} \mid b, isc) &= 0.80 \\
\Pr(CT = \text{tum} \mid b, \overline{isc}) &= 0.90 \\
\Pr(CT = \text{tum} \mid \overline{b}, isc) &= 0.05 \\
\Pr(CT = \text{tum} \mid \overline{b}, \overline{isc}) &= 0.10 \\[4pt]
\Pr(CT = \text{fract} \mid b, isc) &= 0.18 \\
\Pr(CT = \text{fract} \mid b, \overline{isc}) &= 0.01 \\
\Pr(CT = \text{fract} \mid \overline{b}, isc) &= 0.55 \\
\Pr(CT = \text{fract} \mid \overline{b}, \overline{isc}) &= 0.40 \\[4pt]
\Pr(CT = \text{les} \mid b, isc) &= 0.02 \\
\Pr(CT = \text{les} \mid b, \overline{isc}) &= 0.09 \\
\Pr(CT = \text{les} \mid \overline{b}, isc) &= 0.40 \\
\Pr(CT = \text{les} \mid \overline{b}, \overline{isc}) &= 0.50
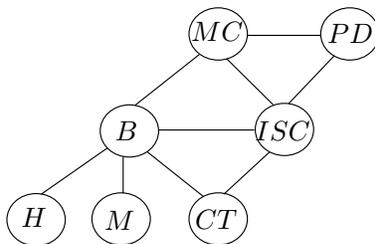\end{aligned}
$$

Figure 1: The *Brain tumor* network

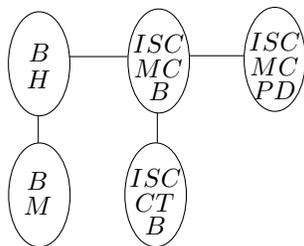Figure 2: The moral graph obtained from the Brain Tumor network

Figure 3: A tree-decomposition of the moralization of the Brain Tumor network

$i \in I\}$ *is a family of subsets (called bags) of* $\mathbf{V}$*, one for each node of* T*, such that*

- $\bigcup_{i \in I} \mathbf{X_i} = \mathbf{V}$,

- *for every edge* $(V, W) \in \mathbf{E}$ *there exists an* $i \in I$ *with* $V \in \mathbf{X_i}$ *and* $W \in \mathbf{X_i}$, *and*

- *for every* $i, j, k \in I$*: if* j *is on the path from* i *to* k *in* T*, then* $\mathbf{X_i} \cap \mathbf{X_k} \subseteq \mathbf{X_j}$.

The width of a tree-decomposition $\langle (I, F), \{\mathbf{X_i} \mid i \in I\} \rangle$ *is* $\max_{i \in I} \mid \mathbf{X_i} \mid -1$.

Treewidth is defined such that a tree (an undirected graph without cycles) has treewidth 1. A polytree (a directed acyclic graph that has no undirected cycles as well) with at most k parents per node has treewidth k. A tree-decomposition of the moralization of the *Brain Tumor* network is shown in Figure 3. The width of this tree-decomposition is 2, since this decomposition has at most 3 variables in each bag.

2.2 Computational Complexity

In the remainder, we assume that the reader is familiar with basic concepts of computational complexity theory, such as Turing Machines, the complexity classes P and NP, and NP-completeness proofs. For more background we refer to classical textbooks like [24] and [25]. In addition to these basic concepts, to describe the complexity of various problems we will use the *probabilistic* class PP, oracles, and fixed-parameter tractability.

The class PP contains languages L accepted in polynomial time by a *Probabilistic Turing Machine*. Such a machine augments the more traditional non-deterministic Turing Machine with a probability distribution associated with each state transition, e.g. by providing the machine with a tape, randomly filled with symbols [28]. If all choice points are binary and the probability of each transition is $\frac{1}{2}$, then the *majority* of the computation paths accept a string s if and only if $s \in L$. This majority, however, is not fixed and may depend on the input, e.g., a problem in PP may accept 'yes'-instances with size n with probability $\frac{1}{2} + \frac{1}{2^n}$. This makes problems in PP intractable in general, in contrast to the related complexity class BPP which is associated with problems which allow for efficient randomized computation. BPP, however, accepts 'yes'-inputs with a *bounded* majority (say $\frac{3}{4}$). This means we can amplify the probability of a correct answer arbitrary close to one by running the algorithm a polynomial amount of times and taking a majority vote on the outcome. This approach failes for unbounded majorities as $\frac{1}{2} + \frac{1}{2^n}$ as allowed by the class PP: here an exponential number of simulations (with respect to the input size) is needed to meet a constant treshhold on the probability.

The canonical PP-complete problem is MAJSAT: given a Boolean formula ϕ, does the majority of the truth instantiations satisfy ϕ? Indeed it is easily shown that MAJSAT encodes SATISFIABILITY: take a formula ϕ with n variables and construct $\psi = \phi \vee x_{n+1}$. Now, the majority of truth assignments satisfy ψ if and only if ϕ is satisfiable, thus NP $\subseteq$ PP. In the field of probabilistic networks, the problem of determining whether the probability $\Pr(\mathbf{X} = \mathbf{x}) \geq q$ (known as the INFERENCE problem) is PP-complete.

A Turing Machine $\mathcal{M}$ has *oracle access* to languages in the class A, denoted as $\mathcal{M}^{\mathrm{A}}$, if it can "query the oracle" in one state transition, i.e., in $\mathcal{O}(1)$. We can regard the oracle as a 'black box' that can answer membership queries in constant time. For example, NP$^{\mathrm{PP}}$ is defined as the class of languages which are decidable in polynomial time on a non-deterministic Turing Machine with access to an oracle deciding problems in PP. Informally, computational problems related to probabilistic networks that are in NP$^{\mathrm{PP}}$ typically combine some sort of *selecting* with *probabilistic inference*. The canonical NP$^{\mathrm{PP}}$-complete satisfiability variant is E-MAJSAT: given a formula ϕ with variable sets $X_1 \ldots X_k$ and $X_{k+1} \ldots X_n$, is there an instantiation to $X_1 \ldots X_k$ such that the majority of the instantiations to $X_{k+1} \ldots X_n$ satisfy ϕ? Likewise, P$^{\mathrm{NP}}$ and P$^{\mathrm{PP}}$ denote classes of languages decidable in polynomial time on a deterministic Turing Machine with access to an oracle for problems in NP and PP, respectively. The canonical satisfiability variants for P$^{\mathrm{NP}}$ and P$^{\mathrm{PP}}$ are LEXSAT and MIDSAT (given ϕ, what

is the lexicographically first, respectively middle, satisfying truth assignment). These classes are associated with finding optimal solutions or enumerating solutions.

In complexity theory, we are often interested in *decision problems*, i.e., problems for which the answer is *yes* or *no*. Well-known complexity classes like P and NP are defined for decision problems and are formalized using Turing Machines. In this paper we will also encounter *function problems*, i.e., problems for which the answer is a function of the input. For example, the problem of determining whether a solution to a 3Sat instance exists, is in NP; the problem of actually *finding* such a solution is in the corresponding function class FNP. Function classes are defined using Turing Transducers, i.e., machines that not only halt in an accepting state on a satisfying input on its input tape, but also return a result on an output tape.

A problem is called *fixed parameter tractable* for a parameter l [26] if it can be solved in time, exponential *only* in l and polynomial in the input size n, i.e. when the running time is $\mathcal{O}(f(l) \cdot n^c)$ for an arbitrary function f and a constant c, independent of n. In practice, this means that problem instances can be solved efficiently, even when the problem is NP-hard in general, if l is known to be small. If an NP-hard problem Π is *fixed parameter tractable* for a parameter l then l is denoted a *source of complexity* of Π [18].

3   Computational Complexity

The problem of finding the most probable explanation for a set of variables in Bayesian networks has been discussed in the literature using many names, like Most Probable Explanation (MPE) [29], Maximum Probability Assignment (MPA) [30], Scenario-Based Explanation [31], (Partial) Abductive Inference or Maximum A Posteriori hypothesis (MAP) [32]. MAP also doubles to denote the set of variables for which an explanation is sought [30]; for this set, also the term *explanation set* is coined [32]. In recent years, more or less consensus is reached to use the terms MPE and Partial MAP to denote the problem with full, respectively partial evidence. We will use the term *explanation set* to denote the set of variables to be explained, and *intermediate nodes* to denote the variables that constitute neither evidence nor the explanation set. The formal definition of the canonical variants of these problems is as follows.

MPE
Instance: A probabilistic network $\mathcal{B} = (\mathbf{G}, \Gamma)$, where $\mathbf{V}$ is partitioned into a set of evidence nodes $\mathbf{E}$ with a joint value assignment $\mathbf{e}$, and an explanation set $\mathbf{M}$.
Output: The most probable joint value assignment $\mathbf{m}$ to the nodes in $\mathbf{M}$ given evidence $\mathbf{e}$, or $\perp$ if $\Pr(\mathbf{m}|\mathbf{e}) = 0$ for every joint value assignment $\mathbf{m}$ to $\mathbf{M}$.

Partial MAP

Instance: A probabilistic network $\mathcal{B} = (\mathbf{G}, \Gamma)$, where $\mathbf{V}$ is partitioned into a set of evidence nodes $\mathbf{E}$ with a joint value assignment $\mathbf{e}$, a set of intermediate nodes $\mathbf{I}$, and an explanation set $\mathbf{M}$.
Output: The most probable joint value assignment $\mathbf{m}$ to the nodes in $\mathbf{M}$ given evidence $\mathbf{e}$, or $\perp$ if $\Pr(\mathbf{m}|\mathbf{e}) = 0$ for every joint value assignment $\mathbf{m}$ to $\mathbf{M}$.

We assume that the problem instance is encoded using a *reasonable* encoding as is customary in computational complexity theory. For example, we expect that numbers are encoded using binary notation (rather than unary), that probabilities are encoded using rational numbers, and that the number of values for each node in the network is bounded. In principle, it is possible to "cheat" on the complexity results by completely discarding the structure in a network $\mathcal{B}$ and encode n stochastic binary variables using a single node with 2^n values that each represent a particular joint value assignment in the original network. The CPT of this node in the thus created network $\mathcal{B}'$ (and thus the input size of the problem) is exponential in the number of variables in the original network, and thus many computational problems will run in time, polynomial in the input size, which of course does not reflect the actual intractability of this approach.

In the next sections we will discuss the complexity of MPE and PARTIAL MAP, respectively. We then enhance both problems to *enumeration* variants: instead of finding the most probable assignment to the explanation set, we are interested in the complexity of finding the k-th most probable assignment for arbitrary values of k. Lastly, we discuss the complexity of *approximating* MPE and PARTIAL MAP and their *parameterized* complexity.

4 MPE and variants

Shimony [33] first addressed the complexity of the MPE problem. He showed that the decision variant of MPE was NP-complete, using a reduction from VERTEX COVER. As already pointed out by Shimony, reductions from several problems are possible, yet using VERTEX COVER allows particular constraints on the structure of the network to be preserved. In particular, it was shown that MPE remains NP-hard, even if all variables are binary and both indegree and outdegree of the nodes is at most two [33].

An alternative proof, using a reduction from SATISFIABILITY, will be given below. In this proof, we need to relax the constraint on the outdegree of the nodes, however, in this variant MPE remains NP-hard when all variables have either uniformly distributed prior probabilities (i.e., $\Pr(V = \text{TRUE}) = \Pr(V = \text{FALSE}) = \frac{1}{2}$) or have deterministic conditional probabilities ($\Pr(V = \text{TRUE} \mid \pi(V))$ is either 0 or 1). The main merit of this alternative proof is, however, that a reduction from SATISFIABILITY may be more familiar for readers not acquainted with graph problems. We first define the decision variant of MPE:
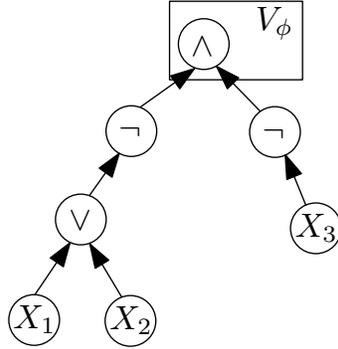
MPE-D

Figure 4: The probabilistic network corresponding to $\neg(X_1 \vee X_2) \wedge \neg X_3$

Instance: A probabilistic network $\mathcal{B} = (\mathbf{G}, \Gamma)$, where $\mathbf{V}$ is partitioned into a set of evidence nodes $\mathbf{E}$ with a joint value assignment $\mathbf{e}$, and an explanation set $\mathbf{M}$; a rational number $0 \leq q < 1$.
Question: Is there a joint value assignment $\mathbf{v}$ to the nodes in $\mathbf{M}$ given evidence $\mathbf{e}$ with probability $\Pr(\mathbf{v} \,|\, \mathbf{e}) > q$?

Let ϕ be a Boolean formula with n variables. We construct a probabilistic network $\mathcal{B}_\phi$ from ϕ as follows. For each propositional variable X_i in ϕ, a binary stochastic variable X_i is added to $\mathcal{B}_\phi$, with possible values TRUE and FALSE and a uniform probability distribution. These variables will be denoted as truth-setting variables $\mathbf{X}$. For each logical operator in ϕ, an additional binary variable in $\mathcal{B}_\phi$ is introduced, whose parents are the variables that correspond to the input of the operator, and whose conditional probability table is equal to the truth table of that operator. For example, the value TRUE of a stochastic variable mimicking the *and*-operator would have a conditional probability of 1 if and only if both its parents have the value TRUE, and 0 otherwise. These variables will be denoted as truth-maintaining variables $\mathbf{T}$. The variable in $\mathbf{T}$ associated with the top-level operator in ϕ is denoted as V_ϕ. The explanation set $\mathbf{M}$ is $\mathbf{V} \setminus V_\phi$. In Figure 4 the network $\mathcal{B}_{\phi_{\mathrm{ex}}}$ is shown for the formula $\phi_{\mathrm{ex}} = \neg(X_1 \vee X_2) \wedge \neg X_3$.

Now, for any particular truth assignment $\mathbf{x}$ to the set of truth-setting variables $\mathbf{X}$ in the formula ϕ we have that the probability of the value TRUE of V_ϕ, given the joint value assignment to the stochastic variables matching that truth assignment, equals 1 if $\mathbf{x}$ satisfies ϕ, and 0 if $\mathbf{x}$ does not satisfy ϕ. With evidence $V_\phi = $ TRUE, the probability of any joint value assignment to $\mathbf{M}$ is 0 if the assignment to $\mathbf{X}$ does not satisfy ϕ, or the assignment to $\mathbf{T}$ does not match the constraints imposed by the operators. However, the probability of any satisfying (and matching) joint value assignment to $\mathbf{M}$ is $\frac{\#_\phi}{2^n}$, where $\#_\phi$ is the number of satisfying truth assignments to ϕ. Thus there exists an instantiation $\mathbf{m}$ to $\mathbf{M}$ such that $\Pr(\mathbf{m} \,|\, V_\phi = $ TRUE$) > 0$ if and only if ϕ is satisfiable. Note that the

9

above network $\mathcal{B}_\phi$ can be constructed from ϕ in polynomial time.

Result 4.1. MPE-D *is* NP-*complete, even when all variables are binary, the indegree of all variables is at most two, and either the outdegree of all variables is two or the probabilities of all variables are deterministic or uniformly distributed.*

Corollary 4.2. MPE *is* NP-*hard under the same constraints as above.*

The exact complexity of the *functional* variant of MPE is discussed in [34]. The proof uses a similar construction as above, however, the prior probabilities of the truth-setting variables is not uniform, but depends on the index of the variable. More in particular, the prior probabilities $p_1, \ldots, p_i, \ldots, p_n$ for the variables $X_1, \ldots, X_i, \ldots, X_n$ are such that $p_i = \frac{1}{2} - \frac{2^i - 1}{2^{n+1}}$. This ensures that a joint value assignment $\mathbf{x}$ to $\mathbf{X}$ is more probable than $\mathbf{x}'$ if and only if the corresponding truth assignment $\mathbf{x}_\phi$ to $X_1, \ldots, X_n$ is *lexicographically* ordered before $\mathbf{x}'_\phi$. Using this construction, Kwisthout [34] reduced MPE from the LexSat-problem of finding the lexicographically first satisfying truth assignment to a formula ϕ. This shows that MPE is $\mathsf{FP}^{\mathsf{NP}}$-complete and thus in the same complexity class as the functional variant of the Traveling Salesman-problem [35].

Result 4.3 ([34]). MPE *is* $\mathsf{FP}^{\mathsf{NP}}$-*complete, even when all variables are binary and the indegree of all variables is at most two.*

Kwisthout [34, p. 70] furthermore argued that the proposed decision variant MPE-D does not capture the essential complexity of the functional problem, and suggested the alternative decision variant MPE-D': given $\mathcal{B}$ and a designated variable $M \in \mathbf{M}$ with designated value m, does M have the value m in the most probable joint value assignment $\mathbf{m}$ to $\mathbf{M}$? This problem turns out to be P^{NP}-complete, using a reduction from the decision variant of LexSat.

Result 4.4 ([34]). MPE-D' *is* P^{NP}-*complete, even when all variables are binary and the indegree of all variables is at most two.*

Bodlaender et al. [30] used a reduction from 3Sat in order to prove a number of complexity results for related problem variants. A 3Sat instance (U, C), where U denotes the variables and C the clauses, was used to construct a probabilistic network $\mathcal{B}_{(U,C)}$ with explanatory set $\mathbf{X} \cup Y$. The construction was such that for any joint value instantiation $\mathbf{x}$ to $\mathbf{X} \cup Y$ that set Y to TRUE, $\mathbf{x}$ was the most probable explanation for $\mathcal{B}_{(U,C)}$ if (U, C) was *not* satisfiable, and the second most probable explanation if if (U, C) *was* satisfiable. Using this construction, they proved (among others) the following complexity results.

Result 4.5 ([30]). *The* IS-AN-MPE *problem (given a network* $\mathcal{B} = (G, \Gamma)$*, an explanatory set* $\mathbf{M}$*, evidence* $\mathbf{e}$*, and an joint value assignment* $\mathbf{m}$ *to* $\mathbf{M}$*: is* $\mathbf{m}$ *the most probable joint value assignment*[1] *to* $\mathbf{M}$*) is* co-NP-*complete.*

---

[1] Or one of the most probable assignments in case of a tie.

Result 4.6 ([30]). *The* BETTER-MPE *problem (given a network* $\mathcal{B} = (G, \Gamma)$, *an explanatory set* $\mathbf{M}$, *evidence* $\mathbf{e}$, *and an joint value assignment* $\mathbf{m}$ *to* $\mathbf{M}$: *find a joint value assignment* $\mathbf{m}'$ *to* $\mathbf{M}$ *which has a higher probability than to* $\mathbf{m}$*) is* NP-*hard.*

5 Partial MAP

Park and Darwiche [36] showed that the decision variant of PARTIAL MAP is NP^{PP}-complete, using a reduction from MAJSAT (given a Boolean formula ϕ partitioned in two sets $\mathbf{X_A}$ and $\mathbf{X_B}$: is there an truth instantiation to $\mathbf{X_A}$ such that the majority of the truth instantiations to $\mathbf{X_B}$ satisfies ϕ?). The proof structure is similar to the hardness proof of MPE, however, the nodes modeling truth setting variables are partitioned into the evidence set $\mathbf{X_A}$ and a set of intermediate variables $\mathbf{X_B}$. Furthermore, q is set to $\frac{1}{2}$. Using this structure NP^{PP}-completeness is proven with the same constraints on the network structure as in Result 4.1. However, Park and Darwiche also prove a considerably strengthened theorem, using an other (and notably more technical) proof:

Result 5.1 ([36]). PARTIAL MAP-*D remains* NP^{PP}-*complete when the network has depth* 2, *there is no evidence, all variables are binary, and all probabilities lie in the interval* $[\frac{1}{2} - \epsilon, \frac{1}{2} + \epsilon]$ *for any fixed* $\epsilon > 0$.

Park and Darwiche [36] show that a number of restricted problem variants remain hard. If there are no intermediate variables, the problem degenerates to MPE-D and thus remains NP-complete. On the other hand, if the explanation set is empty, then the problem degenerates to INFERENCE and thus remains PP-complete. If the number of variables in the explanation set is logarithmic in the total number of variables the problem is in P^{PP} since we can iterate over all joint value assignments of the explanation set in polynomial time and infer the joint probability using an oracle for INFERENCE. If the number of intermediate variables is logarithmic in the total number of variables the problem is in NP since we can verify in polynomial time whether the probability of any given assignment to the variables in the explanation set exceeds the threshold, by summing over the polynomially bounded number of joint value assignments of the other variables. However, when the number of variables in the explanation set or the number of intermediate variables is $\mathcal{O}(n^\epsilon)$ the problem remains NP^{PP}-complete, since we can 'blow up' the general proof construction with a polynomial number of unconnected and deterministic dummy variables such that these constraints are met. Lastly, the problem remains NP-complete when the network is restricted to a polytree.

Result 5.2 ([36]). PARTIAL MAP-*D remains* NP-*complete when restricted to polytrees.*

It follows as a corollary that the functional problem variant PARTIAL MAP is NP^{PP}-hard in general with the same constraints as the decision variant. In addition, Kwisthout [34] shows that PARTIAL MAP is FP$^{NP^{PP}}$-complete and

remains $\mathsf{FP}^{\mathsf{NP}}$-complete on polytrees. This result shares the constraints with Result 4.3.

Result 5.3 ([34]). PARTIAL MAP *is* $\mathsf{FP}^{\mathsf{NP}^{\mathsf{PP}}}$*-complete, even when all variables are binary and the indegree of all variables is at most two.*

Result 5.4 ([34]). PARTIAL MAP *remains* $\mathsf{FP}^{\mathsf{NP}}$*-complete on polytrees, even when all variables are binary and the indegree of all variables is at most two.*

Some variants of PARTIAL MAP can be formulated. For example, in [37] the COND MAP-D problem was defined as follows: Given a probabilistic network $\mathcal{B} = (\mathbf{G}, \Gamma)$, with explanation set $\mathbf{M}$ and designated variable C with designated value c, and a rational number q; is there a joint value assignment $\mathbf{m}$ to $\mathbf{M}$ such that $\Pr(C = c \mid \mathbf{m}) > q$?

It can be easily shown that the hardness proofs of Park and Darwiche [36] for PARTIAL MAP-D can also be applied, with trivial adjustments, to COND MAP-D.

Result 5.5 ([37, 36]). COND MAP-*D is* $\mathsf{NP}^{\mathsf{PP}}$*-complete, even when all variables are binary and the indegree of all variables is at most two.*

Result 5.6. COND MAP-*D remains* NP*-complete on polytrees, even when all variables are binary and the indegree of all variables is at most two.*

6   Enumeration variants

In practical applications, one often wants to find a number of different joint value assignments with a high probability, rather than just the most probable one [38, 39]. For example, in medical applications, one wants to suggest alternative (but also likely) explanations to a set of observations. One might like to prescribe medication that treats a number of plausible (combinations of) diseases, rather than just the most probable combination. It may also be useful to examine the second-best explanation to gain insight in *how good* the best explanation is, relative to other solutions, or how sensitive it is to changes in the parameters of the network [40].

Kwisthout [41] addressed the computational complexity of MPE and PARTIAL MAP when extended to the k-th most probable explanation, for arbitrary values of k. The construction for the hardness proof of KTH MPE is similar to that of Result 4.3, however, the reduction is made from KTHSAT (given a Boolean formula ϕ, what is the lexicographically k-th satisfying truth assignment?) rather than LEXSAT. It is thus shown that KTH MPE is $\mathsf{FP}^{\mathsf{PP}}$-complete and has a P^{PP}-complete decision variant, even if all nodes have indegree at most two. Finding the k-th MPE is thus considerably harder (i.e., complexity-wise) than MPE, and also harder than the PP-complete INFERENCE-problem in Bayesian networks. The computational power of P^{PP} and $\mathsf{FP}^{\mathsf{PP}}$ (and thus the intractability of KTH MPE) is illustrated by Toda's theorem [42] which states that P^{PP} includes the entire Polynomial Hierarchy (PH).

Result 6.1 ([41]). KTH MPE *is* $\mathsf{FP}^{\mathsf{PP}}$*-complete and has a* P^{PP}*-complete decision variant, even if all nodes have indegree at most two.*

The KTH PARTIAL MAP-problem is even harder than that, under usual assumptions[2] in complexity theory. Kwisthout proved [41] that a variant of the problem with *bounds* on the probabilities (BOUNDED KTH PARTIAL MAP) is $\mathsf{FP}^{\mathsf{PP}^{\mathsf{PP}}}$-complete and has a $\mathsf{P}^{\mathsf{PP}^{\mathsf{PP}}}$-complete decision variant, using a reduction from the KTHNUMSAT-problem (given a Boolean formula ϕ whose variables are partitioned in two subsets $\mathbf{X_A}$ and $\mathbf{X_B}$ and an integer l, what is the lexicographically k-th satisfying truth assignment to $\mathbf{X_A}$ such that exactly l truth assignments to $\mathbf{X_B}$ satisfy ϕ?). The proof of Park and Darwiche [36] that shows that PARTIAL MAP remains NP-complete on polytrees, mentioned in the previous section, can be easily modified to reduce KTH PARTIAL MAP on polytrees from the $\mathsf{FP}^{\mathsf{PP}}$-complete problem KTH3SAT [44], hence finding the k-th Partial MAP on polytrees remains $\mathsf{FP}^{\mathsf{PP}}$-complete.

Result 6.2 ([41]). KTH PARTIAL MAP *is* $\mathsf{FP}^{\mathsf{PP}^{\mathsf{PP}}}$*-complete and has a* $\mathsf{P}^{\mathsf{PP}^{\mathsf{PP}}}$*-complete decision variant, even if all nodes have indegree at most two.*

Result 6.3 ([41]). KTH PARTIAL MAP *remains* $\mathsf{FP}^{\mathsf{PP}}$*-complete on polytrees, even if all nodes have indegree at most two.*

7 Approximation Results

While sometimes NP-hard problems can be efficiently approximated in polynomial time (e.g., algorithms exist that find a solution that may not be optimal, but nevertheless is guaranteed to be within a certain bound), no such algorithms exist for the MPE and PARTIAL MAP problems. In fact, Abdelbar and Hedetniemi showed that there can not exist an algorithm that is guaranteed to find a joint value assignment within any fixed bound of the most probable assignment, unless $\mathsf{P} = \mathsf{NP}$ [45]. That does not imply that heuristics play no role in finding assignments; however, if no further restrictions are assumed on the graph structure or probability distribution, no approximation algorithm is *guaranteed* to find a solution (in polynomial time) that has a probability of at least $\frac{1}{r}$ times the probability of the best explanation, for any fixed r.

In fact, it can be easily shown that no algorithm can guarantee *absolute* bounds as well. As we have seen in Section 4, deciding whether there exist a joint value assignment with a probability larger than q is NP-hard for any q larger than 0. Thus, finding a solution which is 'good enough' is NP-hard in general, where 'good enough' may be defined as a ratio of the probability of the best explanation or as an absolute threshold.

Observe that MPE is a special case of PARTIAL MAP, in which the set of intermediate variables $\mathbf{I}$ is empty, and that the intractability of approximating MPE extends to PARTIAL MAP. Furthermore, Park and Darwiche [36] proved

---

[2]To be more precise, the assumptions that the inclusions in the Counting Hierarchy [43] are strict.

that approximating Partial MAP on polytrees within a factor of 2^{n^ϵ} is NP-hard for any fixed $\epsilon, 0 \leq \epsilon < 1$, where n is the size of the problem.

Result 7.1 ([45]). MPE *cannot be approximated within any fixed ratio unless* P = NP.

Result 7.2 ([33]). MPE *cannot be approximated within a fixed bound unless* P = NP.

8   Fixed Parameter Results

In the previous sections we saw that finding the best explanation in a probabilistic network is NP-hard and NP-hard to approximate as well. These intractability results hold *in general*, i.e., when no further constraints are put on the problem instances. However, polynomial-time algorithms are possible for MPE if certain *problem parameters* are known to be small. In this section, we present known results and corollaries that follow from these results. In particular, we discuss the following parameters: probability (Probability-l MPE, Probability-l Partial MAP), treewidth (Treewidth-l MPE, Treewidth-l Partial MAP), and, for Partial MAP, the number of intermediate variables (Intermediate-l Partial MAP). In all of these problems, the input is a probabilistic network and the parameter l as mentioned. Also, for the Partial MAP variants combinations of these parameters will be discussed, in particular probability and treewidth (Probability-l Treewidth-m Partial MAP) and probability and number of intermediate variables (Probability-l Intermediate-m Partial MAP).

Bodlaender et al. [30] presented an algorithm to decide whether the most probable explanation has a probability larger than q, but where q is seen as a fixed parameter rather than part of the input. The algorithm has a running time of $\mathcal{O}(2^{\frac{\log q}{\log 1-q}} \cdot n)$, where n denotes the number of variables. When q is a fixed parameter (and thus assumed constant), this is linear in n; moreover, the running time decreases when q increases, thus for problem instances where the most probable explanation has a high probability, deciding the problem is tractable. The problem is easily enhanced to a functional problem variant where the most probable assignment (rather than TRUE or FALSE) is returned.

Result 8.1 ([30]). Probability-l MPE *is fixed-parameter tractable.*

Corollary 8.2. *Finding the most probable explanation can be done efficiently if the probability of that explanation is high.*

Sy [29] first introduced an algorithm for finding the most probable explanation, based on junction tree techniques, which in multiply connected graphs runs in time, exponential only in the maximum number of node states of the compound variables. Since the size of the compound variables in the junction

tree is equal to the treewidth of the network plus one, this algorithm is expo-nentially only in the treewidth of the network[3]. Hence, if treewidth is seen as a fixed parameter, then the algorithm runs in polynomial time.

Result 8.3 ([29]). TREEWIDTH-L MPE *is fixed-parameter tractable.*

Corollary 8.4. *Finding the most probable explanation can be done efficiently if the treewidth of the network is low.*

Sy's algorithm [29] in fact finds the k most probable explanations (rather than only *the* most probable) and has a running time of $\mathcal{O}(k \cdot n^{|C|})$, where $|C|$ denotes the maximum number of node states of the compound variables. Since k may become exponential in the size of the network this is in general not polynomial, even with low treewidth; however, if k is regarded as parameter then fixed parameter tractability follows as a corollary.

Result 8.5 ([29]). TREEWIDTH-L KTH MPE *is fixed-parameter tractable.*

Corollary 8.6. *Finding the k-th most probable explanation can be done effi-ciently if both k and the treewidth of the network are low.*

When we consider PARTIAL MAP then restricting either the probability or the treewidth is insufficient to render the problem tractable. This latter result follows from the NP-completeness result of Park and Darwiche [36] for PARTIAL MAP restricted to polytrees with at most 2 parents per node, i.e., networks with treewidth at most 2. Furthermore, it is easy to see that deciding PARTIAL MAP includes solving the INFERENCE problem, even if l, the probability of the most probable explanation, is very high. Assume we have a network $\mathcal{B}$ with designated binary variable V. Deciding whether $\Pr(V = \text{TRUE}) > \frac{1}{2}$ is PP-complete in general (see e.g. [34, p.19-21] for a completeness proof, using a reduction from MAJSAT). We now add a binary variable C to our network, with V as its only parent, and probability table $\Pr(C = \text{TRUE} \mid V = \text{TRUE}) = l + \epsilon$ and $\Pr(C = \text{TRUE} \mid V = \text{FALSE}) = l - \epsilon$ for an arbitrary small value ϵ. Now, $\Pr(C = \text{TRUE}) > l$ if and only if $\Pr(V = \text{TRUE}) > \frac{1}{2}$, so determining whether the most probable explanation of C has a probability larger than l boils down to deciding INFERENCE which is PP-complete.

Result 8.7 ([36]). TREEWIDTH-L PARTIAL MAP *is* NP-*complete for $l \geq 2$.*

Result 8.8. PROBABILITY-L PARTIAL MAP *is* PP-*complete independent of the probability l of the most probable explanation.*

However, the algorithm of Bodlaender et al. [30] can be adapted to find par-tial MAPs as well. The algorithm iterates over a topological sort $1, \ldots, i, \ldots, n$ of the nodes of the network. At one point, the algorithm computes $\Pr(V_{i+1} \mid \mathbf{v})$

---

[3]Note that the number of values per variable may be high, thus rendering the algorithm intractable even for networks with low treewidth. However, the conditional probability distri-bution of each variable is part of the problem instance, so even when there are many values per variable, the algorithm still runs in time, polynomial in the input size.

for a particular joint value assignment **v** to $V_1, \ldots, V_i$. In the paper it is concluded that this can be done in polynomial time since all values of $V_1, \ldots, V_i$ are known at iteration step i. To obtain an algorithm for finding partial MAPs, we just skip any iteration step i if V_i is an intermediate variable, and we compute $\Pr(V_{i+1})$ by computing the probability distribution over the 'missing' values V_i. This can be done in polynomial time if either the number of intermediate variables is fixed or the treewidth of the network is fixed. A similar result can be shown for the COND MAP problem variant.

Result 8.9 (adapted from [30]). PROBABILITY-L TREEWIDTH-M PARTIAL MAP *and* PROBABILITY-L INTERMEDIATE-M PARTIAL MAP *are fixed-parameter tractable.*

Corollary 8.10. *Finding the Partial MAP can be done efficiently if both the probability of the most probable explanation is high, and either the treewidth of the network or the number of intermediate variables is low.*

9 Conclusion

Inference of the most probable explanation is hard in general. Approximating the most probable explanation is hard as well. Furthermore, various problem variants, like finding the k-th MPE, finding a better explanation than the one that is given, and finding best explanations when not all evidence is available is hard. Many problems remain hard under severe constraints.

However, this need not to be 'all bad news' for the computational modeler. MPE *is* tractable when the probability of the most probable explanation is high or when the treewidth of the underlying graph is low. PARTIAL MAP *is* tractable when both constraints are met, to name a few examples. The key question for the modeler is: are these constraints plausible with respect to the phenomenon one wants to model? Is it reasonable to suggest that the phenomenon does not occur when the constraints are violated? For example, when cognitive processes like goal inference are modeled as finding the most probable explanation of a set of variables given partial evidence, is it reasonable to suggest that humans have difficulty inferring actions when the probability of the most probable explanation is low, as suggested by [19]?

We do not claim to have answers to such questions. However, the overview of known results in this paper may aid the computational modeler in finding potential sources of intractability. Whether the outcome is received as a blessing (because empirical results may *confirm* those sources of intractability, thus attributing more credibility to the model) or a curse (because empirical results *refute* those sources of intractability, thus providing counterexamples to the model) is beyond our control.

Acknowledgements

References

[1] Lucas, P. J. F., de Bruijn, N., Schurink, K., and Hoepelman, A. (2000) A probabilistic and decision-theoretic approach to the management of infectious disease at the ICU. *Artificial Intelligence in Medicine*, **3**, 251–279.

[2] van der Gaag, L. C., Renooij, S., Witteman, C. L. M., Aleman, B. M. P., and Taal, B. G. (2002) Probabilities for a probabilistic network: a case study in oesophageal cancer. *Artificial Intelligence in Medicine*, **25**, 123–148.

[3] Wasyluk, H., Onisko, A., and Druzdzel, M. J. (2001) Support of diagnosis of liver disorders based on a causal Bayesian network model. *Medical Science Monitor*, **7**, 327–332.

[4] Geenen, P. L., Elbers, A. R. W., van der Gaag, L. C., and van der Loeffen, W. L. A. (2006) Development of a probabilistic network for clinical detection of classical swine fever. *Proceedings of the Eleventh Symposium of the International Society for Veterinary Epidemiology and Economics*, pp. 667–669.

[5] Kennett, R. J., Korb, K. B., and Nicholson, A. E. (2001) Seabreeze prediction using Bayesian networks. In Cheung, D. W.-L., Williams, G. J., and Li, Q. (eds.), *Proceedings of the Fifth Pacific-Asia Conference on Advances in Knowledge Discovery and Data Mining*, pp. 148–153. Springer Verlag, Berlin.

[6] Cofiño, A. S., Cano, R., Sordo, C., and Gutiérrez, J. M. (2002) Bayesian networks for probabilistic weather prediction. In van Harmelen, F. (ed.), *Proceedings of the Fifteenth European Conference on Artificial Intelligence*, pp. 695–699. IOS Press, Amsterdam.

[7] Demirer, R., Mau, R., and Shenoy, C. (2006) Journal of applied finance. *Bayesian Networks: A Decision Tool to Improve Portfolio Risk Analysis*, **16**, 106–119.

[8] Gemela, J. (2001) Financial analysis using Bayesian networks. *Applied Stochastic Models in Business and Industry*, **17**, 57–67.

[9] Kragt, M. E., Newhama, L. T. H., and Jakemana, A. J. (2009) A Bayesian network approach to integrating economic and biophysical modelling. In Anderssen, R. S., Braddock, R. D., and Newham, L. T. H. (eds.), *Proceedings of the 18th World IMACS / MODSIM Congress on Modelling and Simulation*, pp. 2377–2383.

[10] Nedevschi, S., Sandhu, J. S., Pal, J., Fonseca, R., and Toyama, K. (2006) Bayesian networks: an exploratory tool for understanding ICT adoption. *Proceedings of the International Conference on Information and Communication Technologies and Development*, pp. 277–284.

[11] Baker, C. L., Saxe, R., and Tenenbaum, J. B. (2009) Action understanding as inverse planning. *Cognition*, **113**, 329–349.

[12] Yuille, A. and Kersten, D. (2006) Vision as Bayesian inference: analysis by synthesis? *TRENDS in Cognitive Sciences*, **10**, 301–308.

[13] N. Chater, J. B. T. and Yuille, A. (2006) Probabilistic models of cognition: Conceptual foundations. *TRENDS in Cognitive Sciences*, **10**, 287–291.

[14] P.Gärdenfors (1988) *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge, MA.

[15] Pearl, J. (1988) *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, Palo Alto, CA.

[16] Poole, D. and Provan, G. M. (1990) What is the most likely diagnosis? In Bonissone, P., Henrion, M., Kanal, L., and Lemmer, J. (eds.), *Proceedings of the Sixth Annual Conference on Uncertainty in Artificial Intelligence*, pp. 89–106. Elsevier Science, New York, NY.

[17] Chajewska, U. and Halpern, J. (1997) Defining explanation in probabilistic systems. In Geiger, D. and Shenoy, P. (eds.), *Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence*, pp. 62–71. Morgan Kaufmann, San Francisco, CA.

[18] van Rooij, I. and Wareham, T. (2008) Parameterized complexity in cognitive modeling: Foundations, applications and opportunities. *The Computer Journal*, **51**, 385–404.

[19] Blokpoel, M., Kwisthout, J., van der Weide, T. P., and van Rooij, I. (2010) How action understanding can be rational, Bayesian *and* tractable. Manuscript under review for Proceedings of CogSci2010.

[20] Lacave, C. and Díez, F. J. (2002) A review of explanation methods for Bayesian networks. *The Knowledge Engineering Review*, **17**, 107–127.

[21] Jensen, F. V. and Nielsen, T. D. (2007) *Bayesian Networks and Decision Graphs*, second edition. Springer Verlag, New York, NY.

[22] Robertson, N. and Seymour, P. D. (1986) Graph minors II: Algorithmic aspects of tree-width. *Journal of Algorithms*, **7**, 309–322.

[23] Kloks, T. (1994) *Treewidth* LNCS 842. Springer-Verlag, Berlin.

[24] Garey, M. R. and Johnson, D. S. (1979) *Computers and Intractability. A Guide to the Theory of NP-Completeness.* W. H. Freeman and Co., San Francisco, CA.

[25] Papadimitriou, C. H. (1994) *Computational Complexity.* Addison-Wesley.

[26] Downey, R. G. and Fellows, M. R. (1999) *Parameterized complexity.* Springer Verlag, Berlin.

[27] Cooper, G. F. (1984) NESTOR: A computer-based medical diagnostic aid that integrates causal and probabilistic knowledge. Technical Report HPP-84-48. Stanford University.

[28] Gill, J. T. (1977) Computational complexity of Probabilistic Turing Machines. *SIAM Journal of Computing*, **6**.

[29] Sy, B. K. (1992) Reasoning MPE to multiply connected belief networks using message-passing. In Rosenbloom, P. and Szolovits, P. (eds.), *Proceedings of the Tenth National Conference on Artificial Intelligence*, pp. 570–576. AAAI Press, Arlington, Va.

[30] Bodlaender, H. L., van den Eijkhof, F., and van der Gaag, L. C. (2002) On the complexity of the MPA problem in probabilistic networks. In van Harmelen, F. (ed.), *Proceedings of the 15th European Conference on Artificial Intelligence*, pp. 675–679.

[31] Henrion, M. and Druzdzel, M. J. (1990) Qualitative propagation and scenario-based approaches to explanation of probabilistic reasoning. In Bonissone, P., Henrion, M., Kanal, L., and Lemmer, J. (eds.), *Proceedings of the Sixth Conference on Uncertainty in Artificial Intelligence*, pp. 10–20. Elsevier Science, New York, NY.

[32] Neapolitan, R. E. (1990) *Probabilistic Reasoning in Expert Systems. Theory and Algorithms.* Wiley/Interscience, New York, NY.

[33] Shimony, S. E. (1994) Finding MAPs for belief networks is NP-hard. *Artificial Intelligence*, **68**, 399–410.

[34] Kwisthout, J. (2009) The Computational Complexity of Probabilistic Networks. PhD thesis Faculty of Science, Utrecht University, The Netherlands.

[35] Krentel, M. W. (1988) The complexity of optimization problems. *Journal of Computer and System Sciences*, **36**, 490–509.

[36] Park, J. D. and Darwiche, A. (2004) Complexity results and approximation settings for MAP explanations. *Journal of Artificial Intelligence Research*, **21**, 101–133.

[37] Chickering, M. and Halpern, J. (eds.) (2004) *Monotonicity in Bayesian Networks*. Arlington: AUAI press.

[38] Santos, E. (1991) On the generation of alternative explanations with implications for belief revision. In D'Ambrosio, B., Smets, P., and Bonissone, P. (eds.), *Proceedings of the Seventh Conference on Uncertainty in Artificial Intelligence*, pp. 339–347. Morgan Kaufmann, San Mateo, CA.

[39] Charniak, E. and Shimony, S. E. (1994) Cost-based abduction and MAP explanation. *Artificial Intelligence*, **66**, 345–374.

[40] Chan, H. and Darwiche, A. (2006) On the robustness of most probable explanations. *Proceedings of the 22nd Conference on Uncertainty in Artificial Intelligence*, pp. 63–71.

[41] Kwisthout, J. (2008) Complexity results for enumerating MPE and Partial MAP. In Jaeger, M. and Nielsen, T. (eds.), *Proceedings of the Fourth European Workshop on Probabilistic Graphical Models*, pp. 161–168.

[42] Toda, S. (1991) PP is as hard as the polynomial-time hierarchy. *SIAM Journal of Computing*, **20**, 865–877.

[43] Torán, J. (1991) Complexity classes defined by counting quantifiers. *Journal of the ACM*, **38**, 752–773.

[44] Toda, S. (1994) Simple characterizations of P(#P) and complete problems. *Journal of Computer and System Sciences*, **49**, 1–17.

[45] Abdelbar, A. M. and Hedetniemi, S. M. (1998) Approximating MAPs for belief networks is NP-hard and other theorems. *Artificial Intelligence*, **102**, 21–38.