Most Probable Explanations In Bayesian Networks: Complexity And Tractability

Johan Kwisthout

Radboud University Nijmegen
Institute for Computing and Information Sciences
P.O. Box 9010, 6500GL Nijmegen, The Netherlands.

Abstract

One of the key computational problems in Bayesian networks is computing the maximal posterior probability of a set of variables in the network, given an observation of the values of another set of variables. In its most simple form, this problem is known as the MPE-problem. In this paper, we give an overview of the computational complexity of many problem variants, including enumeration variants, parameterized problems, and approximation strategies to the MPE-problem with and without additional (neither observed nor explained) variables. Many of these complexity results appear elsewhere in the literature; other results have not been published yet. The paper aims to provide a fairly exhaustive overview of both the known and new results.

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 diagnosis, like the systems described in [1–4], 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 [7–9], sociology [10,11], and cognitive tasks as vision or goal inference [12,13]. Although these tasks may superficially appear different, the underlying computational problem is the same: given a probabilistic

Email address: `johank@science.ru.nl` (Johan Kwisthout).

network, describing a set of stochastic variables and the conditional independencies 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 [14]), 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 [15], [16], [17], and [18]). 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 sometimes found in the literature (e.g. in [14]) 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 [19]. However, when constraints are imposed on the structure of the network or on the probability distribution, the problem may become tractable. In other words: the optimization criterion is not a *source of complexity*[20] 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.

The paper is intended 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 [21]. Here the Bayesian Inverse Planning model [12], a cognitive model for human goal inference based on Bayesian abduction, was studied and—based on computational complexity analysis—the conditions under which the model becomes intractable, respectively remains tractable were identified, allowing the modelers to investigate the (psychological) plausibility of these conditions. For example, using complexity analysis they concluded that the model predicts that if people have many parallel goals that influence their actions, it is in general hard for an observer to infer the most probable combination of goals, based on the observed actions; however, if the probability of the most probable combination of goals is high, then inference is tractable again.

While good introductions to explanation problems in Bayesian networks exist (see, e.g., [22] 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. Being aware of these issues (i.e., the constraints that render explanation problems tractable, respectively leave the problems intractable) is in our opinion key to a thorough understanding of the phenomena that are studied [20]. Furthermore, it allows investigators to not only constrain their computational models to be tractable under circumstances where empirical results suggest that the task at hand is tractable indeed, but also to let their models *predict* under which circumstances the task becomes intractable and thus assist in generating hypotheses which may be empirically testable.

In this paper we 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 [16,23–29].

An overview paper on complexity results necessarily contains many complexity classes and computational problems. All complexity classes in this paper that are introduced informally in the main text will be formally defined in Appendix A. For easy reference, all computational problems are also formally defined in Appendix B.

A Bayesian or probabilistic network $\mathcal{B}$ is a graphical structure that models a set of stochastic variables, the conditional independencies among these variables, and a joint probability distribution over these variables. $\mathcal{B}$ includes a directed acyclic graph $\mathbf{G}_{\mathcal{B}} = (\mathbf{V}, \mathbf{A})$, modeling the variables and conditional independencies 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 $\Pr(\mathbf{V}) = \prod_{i=1}^{n} \Pr(V_i \mid \pi(V_i))$ over its variables, where $\pi(V_i)$ denotes the parents of V_i in $\mathbf{G}_{\mathcal{B}}$. 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 $\mathbf{e}$ is observed. We will sometimes write $\Pr(\mathbf{h} \mid \mathbf{e})$ as a shorthand for $\Pr(\mathbf{H} = \mathbf{h} \mid \mathbf{E} = \mathbf{e})$ if no ambiguity can occur.

A small example of a Bayesian network is the *Brain Tumor* network, shown in Figure 1. This network, adapted from Cooper [30], 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 also impaired. Furthermore, it is likely that a CT-scan (CT) of the head will reveal a tumor if it is present.

Every (posterior) probability of interest in Bayesian networks can be computed using well known lemmas in probability theory, like Bayes' theorem ($\Pr(H \mid E) = \frac{\Pr(E \mid H) \times \Pr(H)}{\Pr(E)}$), marginalization ($\Pr(H) = \sum_{g_i} \Pr(H \wedge G = g_i)$), and the factorization property ($\Pr(\mathbf{V}) = \prod_{i=1}^{n} \Pr(V_i \mid \pi(V_i))$) of Bayesian networks. For example, from the definition of the Brain Tumor network we can compute that $\Pr(B = \textsc{true} \mid M = \textsc{true}, CT = \textsc{false}) = 0.09$ and that $\Pr(MC = \textsc{true}, PD = \textsc{false} \mid M = \textsc{false}, H = \text{absent}) = 0.13$.

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: when the variables tend to be clustered in small groups with few connections between groups, treewidth is typically low, whereas treewidth tends to be high if there are no clear clusters and the connections between variables are scattered all over the network. Treewidth plays an important role in the complexity analysis of Bayesian networks, as many otherwise intractable computational problems

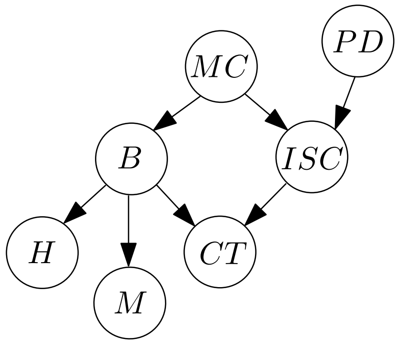

$\Pr(MC = \textsc{true}) = 0.20$

$\Pr(PD = \textsc{true}) = 0.10$

$\Pr(B = \textsc{true} \mid MC = \textsc{true}) = 0.20$

$\Pr(B = \textsc{true} \mid MC = \textsc{false}) = 0.05$

$\Pr(ISC = \textsc{true} \mid MC = \textsc{true}, PD = \textsc{true}) = 0.95$

$\Pr(ISC = \textsc{true} \mid MC = \textsc{true}, PD = \textsc{false}) = 0.80$

$\Pr(ISC = \textsc{true} \mid MC = \textsc{false}, PD = \textsc{true}) = 0.70$

$\Pr(ISC = \textsc{true} \mid MC = \textsc{false}, PD = \textsc{false}) = 0.20$

$\Pr(H = \text{severe} \mid B = \textsc{true}) = 0.70$

$\Pr(H = \text{moderate} \mid B = \textsc{true}) = 0.25$

$\Pr(H = \text{absent} \mid B = \textsc{true}) = 0.05$

$\Pr(H = \text{severe} \mid B = \textsc{false}) = 0.30$

$\Pr(H = \text{moderate} \mid B = \textsc{false}) = 0.20$

$\Pr(H = \text{absent} \mid B = \textsc{false}) = 0.50$

$\Pr(M = \textsc{true} \mid B = \textsc{true}) = 0.50$

$\Pr(M = \textsc{true} \mid B = \textsc{false}) = 0.30$

$\Pr(CT = \textsc{true} \mid B = \textsc{true}, ISC = \textsc{true}) = 0.90$

$\Pr(CT = \textsc{true} \mid B = \textsc{true}, ISC = \textsc{false}) = 0.80$

$\Pr(CT = \textsc{true} \mid B = \textsc{false}, ISC = \textsc{true}) = 0.10$

$\Pr(CT = \textsc{true} \mid B = \textsc{false}, ISC = \textsc{false}) = 0.05$

Fig. 1. The *Brain Tumor* network with its conditional probability distributions

become tractable when the treewidth of the network is bounded.

We define the treewidth of a Bayesian network $\mathcal{B}$ as the treewidth of a triangulation of the moralization $\mathbf{G}^{\mathbf{M}}_{\mathcal{B}}$ of its graph $\mathbf{G}_{\mathcal{B}}$. This moralization is the undirected graph that is obtained from $\mathbf{G}_{\mathcal{B}}$ by adding arcs so as to connect all pairs of parents of a variable, and then dropping all directions; we will use the phrase 'moralized graph' to refer to the moralization of the graph of a network. The moralized graph of the *Brain Tumor* network is shown in Figure 2. A triangulation of the moralized graph $\mathbf{G}^{\mathbf{M}}_{\mathcal{B}}$ is any graph $\mathbf{G_T}$ that embeds $\mathbf{G}^{\mathbf{M}}_{\mathcal{B}}$ as a subgraph and in addition is chordal, that is, it does not include loops of more than three variables without any pair being adjacent in $\mathbf{G_T}$. A tree-decomposition [25] of a triangulation $\mathbf{G_T}$ is a tree $\mathbf{T_G}$ such that

- each node $\mathbf{X}_i$ in $\mathbf{T_G}$ is a bag of nodes which constitute a clique in $\mathbf{G_T}$;
- for every i, j, k, if $\mathbf{X}_j$ lies on the path from $\mathbf{X}_i$ to $\mathbf{X}_k$ in $\mathbf{T_G}$, then $\mathbf{X}_i \cap \mathbf{X}_k \subseteq \mathbf{X}_j$.

The width of the tree-decomposition $\mathbf{T_G}$ of the graph $\mathbf{G_T}$ equals $\max_i(|\mathbf{X}_i| - 1)$, that is, it equals the size of the largest clique in $\mathbf{G_T}$, minus 1. The treewidth of a Bayesian network $\mathcal{B}$ now is the minimum width over all possible tree-decompositions of triangulations of $\mathbf{G}^{\mathbf{M}}_{\mathcal{B}}$.

Treewidth is defined such that a tree (an undirected graph without cycles)

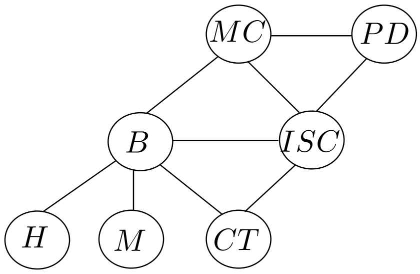

Fig. 2. The moralized graph obtained from the Brain Tumor network

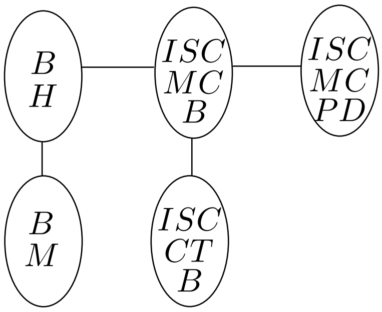

Fig. 3. A tree-decomposition of the moralization of the Brain Tumor network

has treewidth 1. 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. Note that each undirected graph has many tree-decompositions, that may vary in width; recall that the treewidth of such a graph is defined as the *minimal* width over all possible tree-decompositions.

2.2 Computational Complexity Theory

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 [27] and [28]. In addition to these basic concepts, to describe the complexity of various problems we will use the *probabilistic* class PP, oracles, function classes, and some aspects from parameterized complexity theory.

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 [31]. Acceptance of an input x is defined as follows: the probability of arriving in an *accept state* is strictly larger than $\frac{1}{2}$ if and only if $x \in L$. If all choice points are binary and the probability of each transition is $\frac{1}{2}$, then an identical definition is that the *majority* of the computation paths accept an input x if and only if $x \in L$. This probability of acceptance, however, is not fixed and may (exponentially) depend on the input, e.g., a problem in PP may accept 'yes'-instances with size $|x|$ with probability $\frac{1}{2} + \frac{1}{2^{|x|}}$. This potentially small majority makes problems in PP intractable in general: we cannot amplify the probability of acceptance by running a probabilistic algorithm multiple times and taking a majority vote on the output, unless we're prepared to run the algorithm exponentially many times. But in that case, we might as well use brute force to solve the problem exactly.

The canonical PP-complete problem is MAJSAT: given a Boolean formula ϕ, does the majority of the truth assignments satisfy ϕ? Indeed it is easily shown that MAJSAT encodes the NP-complete SATISFIABILITY problem: take a formula ϕ with n variables and construct $\psi = \phi \vee x_{n+1}$. Now, the majority of truth assignments to $x_1 \ldots x_{n+1}$ satisfy ψ if and only if ϕ is satisfiable, thus $\mathsf{NP} \subseteq \mathsf{PP}$. In Bayesian networks, the canonical problem of determining whether the probability $\Pr(\mathbf{H} = \mathbf{h} \mid \mathbf{E} = \mathbf{e}) > q$ for a given rational q (known as the INFERENCE problem) is PP-complete [32].

A Turing Machine $\mathcal{M}$ has *oracle access* to languages in the class A, denoted as $\mathcal{M}^{\mathsf{A}}$, 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, $\mathsf{NP}^{\mathsf{PP}}$ 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 Bayesian networks that are in $\mathsf{NP}^{\mathsf{PP}}$, like PARAMETER TUNING[33], typically combine some sort of *selecting* with *probabilistic inference.* The canonical $\mathsf{NP}^{\mathsf{PP}}$-complete satisfiability variant is E-MAJSAT[34]: given a formula ϕ with variable sets $x_1 \ldots\ x_k$ and $x_{k+1} \ldots\ x_n$, is there an truth assignment to $x_1 \ldots\ x_k$ such that the majority of the truth assignments to $x_{k+1} \ldots\ x_n$ satisfy ϕ? Likewise, P^{NP} and P^{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^{NP} and P^{PP} are LEXSAT and KTHSAT (given ϕ, what is the lexicographically first, respectively k-th, satisfying truth assignment?); P^{NP} and P^{PP} are associated with finding optimal solutions or enumerating solutions, respectively [35,36].

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. There is no one-to-one mapping between decision classes and function classes, in the sense that if a decision problem is in NP, its functional variant is not [1] always in FNP. Common classes for functional variants of NP-complete problems are also $\mathsf{FP}^{\mathsf{NP}}$ (solvable with a deterministic Turing Transducer with access to an oracle for problems in NP) and $\mathsf{FP}^{\mathsf{NP}}_{[\log]}$ (solvable with a deterministic Turing Transducer with *limited* access to an oracle for problems in NP, i.e., at most a logarithmic amount of calls, with respect to the input size). Whether the problem is in FNP, $\mathsf{FP}^{\mathsf{NP}}$ or $\mathsf{FP}^{\mathsf{NP}}_{[\log]}$ depends on the nature of the problem: "$\exists$"-like problems, like 3SAT, are typically in FNP; optimization problems, like TRAVELLING SALESMAN PROBLEM or HAMILTONIAN CIRCUIT are in $\mathsf{FP}^{\mathsf{NP}}$ or $\mathsf{FP}^{\mathsf{NP}}_{[\log]}$, depending on whether the problem definition does (TRAVELLING SALESMAN PROBLEM) or doesn't (HAMILTONIAN CIRCUIT) include *weights* [37].

Sometimes problems are intractable (i.e., NP-hard) in general, but become tractable if some *parameter* of the problem can be assumed to be small. A *parameterized problem* is a pair (Π, κ) of a decision problem Π and a polynomial time computable *parameterization* $\kappa : \{0,1\}^* \to \mathbb{N}$ mapping strings to natural numbers. The parameterized problem (Π, κ) is *fixed-parameter tractable* if there exists an algorithm deciding every instance (x, l) of (Π, κ) with running time $\mathcal{O}(f(\kappa(x, l)) \cdot |x|^c)$ for an arbitrary computable function f and a constant c, independent of $|x|$ [38,29]. The class of all fixed-parameter tractable decision problems is denoted as FPT. To improve readability, if the parameterization is clear from the context (e.g., $\kappa(x, l) = l$), we just mention the parameter l.

Informally, a problem is called fixed-parameter tractable for a parameter l if it can be solved in time, exponential *only* in l and polynomial in the input size $|x|$. 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* [20] of Π: bounding l renders the problem tractable, whereas leaving l unbounded ensures intractability under usual complexity-theoretic assumptions like $\mathsf{P} \neq \mathsf{NP}$. On the other hand, if Π remains NP-hard for all but finitely many values of the parameter l, then Π is para-NP-hard: bounding l does not render the problem tractable. The notion of fixed-parameter tractability can be extended to deal with *rational*, rather

[1] Given the usual assumptions in computational complexity theory; in this case, that $\mathsf{FNP} \subsetneq \mathsf{FP}^{\mathsf{NP}}_{[\log]} \subsetneq \mathsf{FP}^{\mathsf{NP}}$.

than integer, parameters [2]. Informally, if a problem is fixed-rational tractable for a (rational) parameter l, then the problem can be solved tractably if l is close to 0. For readability, we will liberally mix integer and rational parameters in the remainder.

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) [39], Maximum Probability Assignment (MPA) [40], Belief Revision [16], Scenario-Based Explanation [41], Marginal MAP [24], (Partial) Abductive Inference or Maximum A Posteriori hypothesis (MAP) [42]. MAP also doubles to denote the set of variables for which an explanation is sought [40]; for this set, also the term *explanation set* is coined [42]. 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.

For example, in the Brain Tumor network one could be interested in the most probable joint value assignment to MC (presence of metastatic cancer) and PD (presence of Paget's disease) given the evidence that the patient has severe headaches (H = severe), memory is impaired (M = TRUE) but no tumor could be found on the CT-scan (CT = FALSE). Here, the explanation set is $\{MC, HD\}$; the evidence set is $\{H, M\}$ and the intermediate nodes are $\{B, ISC\}$. In fact, solving the PARTIAL MAP problem with the above specifications would reveal that the most probable joint value assignment would be the *absence* of both metastatic cancer and Paget's disease as this joint value assignment has a conditional probability of 0.78. If we'd also observe the absence of a brain tumor (B = FALSE) yet an increased serum calcium (ISC = TRUE) then the set of intermediate nodes would be empty, and then solving an MPE problem would reveal that PD = TRUE, MC = FALSE would be the most probable joint value assignment, with probability 0.34.

[2] Here a *rational parameterization* is a function $\lambda : \{0,1\}^* \to [0,1\rangle$; a rationally parameterized problem (Π, λ) is *fixed-rational tractable* if there exists an algorithm deciding every instance (x, l) of (Π, λ) with running time $\mathcal{O}(f(\lambda(x,l)) \cdot |x|^c)$ for a *nondecreasing* computable function $f : \mathbb{Q} \cap [0,1\rangle$ and a constant c, independent of $|x|$. Every fixed-rational problem (Π, λ) can be translated in a fixed-parameter problem $(\Pi, \lambda)^t$ using a translation $t : \mathbb{N} \to [0,1\rangle$ (Moritz M. Müller, personal communication).

The formal definition of the canonical variants of the MPE and PARTIAL MAP problems is as follows.

MPE
Instance: A probabilistic network $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Gamma)$, 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: $\arg\max_{\mathbf{m}} \Pr(\mathbf{m}, \mathbf{e})$, i.e., the most probable joint value assignment $\mathbf{m}$ to the nodes in $\mathbf{M}$ and 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 $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Gamma)$, 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: $\arg\max_{\mathbf{m}} \Pr(\mathbf{m}, \mathbf{e})$, i.e., the most probable joint value assignment $\mathbf{m}$ to the nodes in $\mathbf{M}$ and evidence $\mathbf{e}$, or $\perp$ if $\Pr(\mathbf{m}, \mathbf{e}) = 0$ for every joint value assignment $\mathbf{m}$ to $\mathbf{M}$.

Note that in the above definition we use $\arg\max_{\mathbf{m}} \Pr(\mathbf{m}, \mathbf{e})$, rather than $\arg\max_{\mathbf{m}} \Pr(\mathbf{m} \mid \mathbf{e})$ which is probably more often seen in the literature. Observe, however, that $\Pr(\mathbf{m} \mid \mathbf{e}) = \frac{\Pr(\mathbf{m},\mathbf{e})}{\Pr(\mathbf{e})}$, and the *value* of the most probable joint value assignment (i.e., not its *probability*) is independent of $\Pr(\mathbf{e})$ as this is equal for every joint value assignment. However, when we examine the decision variant of MPE, we will see that there is a difference in computational complexity between the "joint" and the "conditional" notion of MPE. We will denote the latter problem (i.e., find the *conditional* MPE or $\arg\max_{\mathbf{m}} \Pr(\mathbf{m} \mid \mathbf{e})$) as MPEE in line with [43]. A similar variant exists for the PARTIAL MAP-problem, however we will argue that the computational complexity of these problems is identical and we will use both problems variants liberally in further results.

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 variable in the network is bounded by a constant, unless explicitly mentioned otherwise. In principle, it is possible to "cheat" on the complexity results by completely discarding the structure (i.e., the independency relations) 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 variants 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 [44] 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. While reductions from several problems are possible, the use of Vertex Cover allowed 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 [44]. The intractability of MPE is due to the fact that we may need to consider an exponential number of joint value assignments.

An alternative proof, using a reduction from Satisfiability, will be given below. In this proof (in its original form originating from [45]), 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
Instance: A probabilistic network $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Gamma)$, 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 $\mathbf{m}$ to the nodes in $\mathbf{M}$ with evidence $\mathbf{e}$ with probability $\Pr(\mathbf{m}, \mathbf{e}) > q$?

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

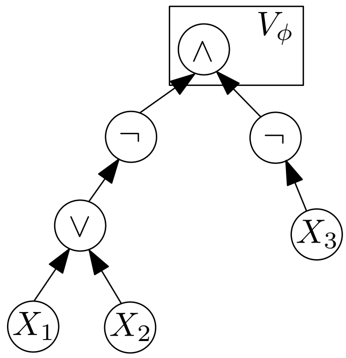

Fig. 4. The probabilistic network corresponding to $\neg(x_1 \vee x_2) \wedge \neg x_3$

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 $\mathbf{X} \cup \mathbf{T} \setminus \{V_\phi\}$. In Figure 4 the network $\mathcal{B}_{\phi_{\text{ex}}}$ is shown for the formula $\phi_{\text{ex}} = \neg(x_1 \vee x_2) \wedge \neg x_3$.

Theorem 1 MPE-D *is* NP-*complete*

Proof. We can prove membership in NP using a certificate consisting of a joint value assignment $\mathbf{m}$. As $\mathcal{B}$ is partitioned into $\mathbf{M}$ and $\mathbf{E}$, we can compute any probability of interest in polynomial time as we have a value assignment for all variables.

To prove hardness, we apply the construction as illustrated above. 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 $V_\phi = \text{TRUE}$, the probability of any joint value assignment to $\mathbf{M}$ is 0 if the assignment to $\mathbf{X}$ does not satisfy ϕ, or if 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 $\frac{1}{\#_\phi}$, where $\#_\phi$ is the number of satisfying truth assignments to ϕ. Thus there exists an joint value assignment $\mathbf{m}$ to $\mathbf{M}$ such that $\Pr(\mathbf{m}, V_\phi = \text{TRUE}) > 0$ if and only if ϕ is satisfiable. Note that the above network $\mathcal{B}_\phi$ can be constructed from ϕ in time, polynomial in the size of ϕ, since we introduce only a single variable for each variable and for each operator in ϕ. $\square$

Result 2 MPE-D *is* NP*-complete, even when all variables are binary, the indegree of all variables is at most two, there are no arcs from the evidence set to the explanation set, and either the outdegree of all variables is two or the probabilities of all variables are deterministic or uniformly distributed.*

Corollary 3 MPE *is* NP*-hard under the same constraints as above.*

The decision variant of the MPEe problem (given a network $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Gamma)$, an explanatory set $\mathbf{M}$, and an evidence set $\mathbf{E}$ with evidence $\mathbf{e}$, compute $\arg\max_{\mathbf{m}} \Pr(\mathbf{m} \mid \mathbf{e})$) was proven PP-complete in [43] by a reduction from Maj3Sat (i.e., Majsat restricted to formulas in 3CNF form). The source of this increase in complexity[3] is the division by $\Pr(\mathbf{e})$ to obtain $\Pr(\mathbf{m} \mid \mathbf{e}) = \frac{\Pr(\mathbf{m},\mathbf{e})}{\Pr(\mathbf{e})}$. Since the set of vertices $\mathbf{V}$ is partitioned into $\mathbf{M}$ and $\mathbf{E}$, computing $\Pr(\mathbf{e})$ is a inference problem which has a PP-complete decision variant.

Result 4 ([43]) MPEe-D *is* PP*-complete, even when all variables are binary.*

The exact complexity of the *functional* variant of MPE is discussed in [46]. The proof uses a similar construction as above, however, the prior probabilities of the truth-setting variables are not uniform, but depend 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 defined as $p_i = \frac{1}{2} - \frac{2^i - 1}{2^{n+1}}$. 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 $\mathbf{x}$ to $x_1, \ldots, x_n$ is *lexicographically* ordered before $\mathbf{x}'$. Using this construction, Kwisthout et al. [46] 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 Travelling Salesman Problem [35].

This result reflects the fact that, like TSP, MPE really is an optimization problem: we are not merely interested in *a* solution that exceeds a threshold probability, but in the *best* solution. While it is easy to verify that a given solution exceeds a threshold (we can compute $\Pr(\mathbf{m}, \mathbf{e})$ in polynomial time given $\mathbf{m}$ and $\mathbf{e}$), there is no apparent trivial way to verify that there is no *better* solution indeed. However, we may use some sort of binary search to find the best solution, but this may take a polynomial number of queries to an oracle solving MPE-D. This is reflected in the $\mathsf{FP}^{\mathsf{NP}}$-completeness of MPE.

Result 5 ([46]) MPE *is* $\mathsf{FP}^{\mathsf{NP}}$*-complete, even when all variables are binary and the indegree of all variables is at most two.*

Kwisthout [33, p. 70] furthermore argued that the proposed decision variant

[3] Under the usual assumption that $\mathsf{NP} \neq \mathsf{PP}$.

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 similar reduction as above, yet now from the *decision* variant of LexSat.

Result 6 ([33]) 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. [40] 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 $\mathcal{B}_{(U,C)}$ as follows. For each variable $x_i \in U$, a binary variable X_i with uniform distribution was added to $\mathcal{B}_{(U,C)}$. In addition, a binary and uniformly distributed variable Y was added. For each clause $c_j \in C$, a binary variable C_j was added to $\mathcal{B}_{(U,C)}$, with the variables from U appearing in c_j and the variable Y as parents. Lastly, a binary variable D was added with Y as parent. The conditional probability $\Pr(C_j \mid \pi(C_j))$ was $\frac{3}{4}$ if Y had the value TRUE or if the corresponding truth assignment to $\pi(C_j) \setminus Y$ satisfied the clause c_j, and $\frac{1}{2}$ otherwise. The conditional probability $\Pr(D \mid Y)$ was $\frac{1}{2}$ if Y was set to TRUE and $\frac{3}{4}$ otherwise.

The construction ensures that for any joint value assignment $\mathbf{x}$ to $X_i \ldots X_n \cup Y$ 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 (U, C) *was* satisfiable. Using this construction, they proved (among others) the following complexity results.

Result 7 ([40]) *The* IS-AN-MPE *problem (given a network* $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Gamma)$*, an explanatory set* $\mathbf{M}$*, evidence set* $\mathbf{E}$ *with evidence* $\mathbf{e}$*, and an joint value assignment* $\mathbf{m}$ *to* $\mathbf{M}$*: is* $\mathbf{m}$ *the most probable joint value assignment*[4] *to* $\mathbf{M}$*) is* co-NP*-complete.*

Result 8 ([40]) *The* BETTER-MPE *problem (given a network* $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Gamma)$*, an explanatory set* $\mathbf{M}$*, evidence set* $\mathbf{E}$ *with 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* $\mathbf{m}$*) is* NP*-hard.*

Intuitively, the IS-AN-MPE is co-NP-hard because if we can decide in polynomial time whether a particular assignment $\mathbf{x}$, with Y set to TRUE, is the most probable explanation of $X_i \ldots X_n \cup Y$, then we can also decide that (U, C) is *not* satisfiable. Similarly, as any assignment $\mathbf{x}$ with Y set to TRUE is the second-best assignment to $X_i \ldots X_n \cup Y$ if and only if (U, C) is satisfiable, if

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

we can decide whether there is a better explanation then we can also decide (U, C), hence BETTER-MPE is NP-hard.

Lastly, to facilitate a later proof, we define (a decision variant of) the MinPE problem as follows: given a network $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Gamma)$, an explanatory set $\mathbf{M}$, evidence set $\mathbf{E}$ with evidence $\mathbf{e}$ and a rational number q: does $\Pr(\mathbf{m_i}, \mathbf{e}) > q$ hold *for all* joint value assignments $\mathbf{m_i}$ to $\mathbf{M}$? It can be readily seen that this problem is co-NP-complete: membership in co-NP follows since we can falsify the claim using a certificate consisting of a suitable joint value assignment $\mathbf{m_i}$ in polynomial time. Hardness can be shown using a similar reduction as used to prove NP-hardness of MPE-D, but now from the canonical co-NP-complete problem Tautology; intuitively, the MinPE problem is hard as there are potentially exponentially many joint value assignments to $\mathbf{M}$ and we must verify that for *all* of them $\Pr(\mathbf{m_i}, \mathbf{e}) > q$ holds.

Result 9 *The* MinPE-D *problem is* co-NP*-complete.*

5 Partial MAP

Park and Darwiche [45] first addressed the computational complexity of Partial MAP. They showed that the decision variant of Partial MAP is $\mathsf{NP}^{\mathsf{PP}}$-complete, using a reduction from E-Majsat (given a Boolean formula ϕ partitioned in two sets $\mathbf{X_E}$ and $\mathbf{X_M}$: is there an truth assignment to $\mathbf{X_E}$ such that the majority of the truth assignments to $\mathbf{X_M}$ satisfies ϕ?). The proof structure is similar to the hardness proof of MPE in Section 4, however, the nodes modeling truth setting variables are partitioned into the evidence set $\mathbf{X_E}$ and a set of intermediate variables $\mathbf{X_M}$. Using this structure $\mathsf{NP}^{\mathsf{PP}}$-completeness is proven with the same constraints on the network structure as in Result 2. However, Park and Darwiche also prove a considerably strengthened theorem, using an other (and notably more technical) proof:

Result 10 ([45]) Partial MAP-D *remains* $\mathsf{NP}^{\mathsf{PP}}$*-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* $\epsilon > 0$.

These complexity results can be intuitively understood when we envisage that for solving Partial MAP there are *two* sources of complexity. One both has to choose a joint value assignment out of potentially exponentially many such assignments (the "NP-part") and, for each such assignment, marginalize over the (also potentially exponentially many) joint value assignments to the intermediate variables (the "PP-part"). However, since we already need the power of the PP-oracle to compute $\Pr(\mathbf{m}, \mathbf{e}) = \sum_{\mathbf{i}} \Pr(\mathbf{m}, \mathbf{e}, \mathbf{I} = \mathbf{i})$, having to compute $\Pr(\mathbf{e})$ to obtain $\Pr(\mathbf{m} \mid \mathbf{e})$ 'does not hurt us' complexity-wise; both the "marginal" and the "conditional" decision variants of Partial MAP are in $\mathsf{NP}^{\mathsf{PP}}$.

Park and Darwiche [45] 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. As we then need to marginalize only over a polynomially bounded number of joint value assignments of the intermediate variables, we can verify in polynomial time whether the probability of any given joint value assignment to the variables in the explanation set exceeds the threshold.

However, when the number of variables in the explanation set or the number of intermediate variables is $\mathcal{O}(n^\epsilon)$ the problem remains $\mathsf{NP}^{\mathsf{PP}}$-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, as shown by Park and Darwiche using a reduction from Maxsat.

Result 11 ([45]) Partial MAP-D *remains* NP*-complete when restricted to polytrees.*

It follows as a corollary that the functional problem variant Partial MAP is $\mathsf{NP}^{\mathsf{PP}}$-hard in general with the same constraints as the decision variant. The exact complexity of the functional variant is discussed in [46]. Using a similar proof construct as for the functional variant of MPE, it was proven that Partial MAP is $\mathsf{FP}^{\mathsf{NP}^{\mathsf{PP}}}$-complete, and that this result shares the constraints with Result 5. The intuition behind this complexity class is also similar to the case of MPE, yet we also need to marginalize over the intermediate variables. This leaves us with *three* aspects of intractability that work on top of each other: finding a candidate solution that exceeds a threshold (the "NP-part"), marginalizing over the intermediate variables (the "PP-part"), and finally using binary search to compute the optimal solution (the "FP-part").

Result 12 ([33]) 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.*

Some variants of Partial MAP have been formulated in the literature. For example, in [47] the CondMAP-D problem was defined as follows: Given a network $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Gamma)$, with explanation set $\mathbf{M}$, evidence set $\mathbf{E}$ with evidence

$\mathbf{e}$, and a rational number q; is there a joint value assignment $\mathbf{m}$ to $\mathbf{M}$ such that $\Pr(\mathbf{e} \mid \mathbf{m}) > q$? Note that CondMAP-D differs from Partial MAP-D as we use $\Pr(\mathbf{e} \mid \mathbf{m})$ rather than $\Pr(\mathbf{m} \mid \mathbf{e})$. In [47] it has been shown that the hardness proof of Park and Darwiche [45] for Partial MAP-D can also be applied, with trivial adjustments, to CondMAP-D, since in their proof construct $\mathbf{M}$ does not have incoming arcs and has a uniform prior distribution for each joint value assignment $\mathbf{m}$. Since $\Pr(\mathbf{e} \mid \mathbf{m}) = \frac{\Pr(\mathbf{e},\mathbf{m})}{\Pr(\mathbf{m})}$ it follows that $\Pr(\mathbf{e} \mid \mathbf{m}) > q$ iff. $\Pr(\mathbf{e}, \mathbf{m}) > \Pr(\mathbf{m}) \times q$.

Result 13 ([47,45]) CondMAP-D *is* $\mathsf{NP}^{\mathsf{PP}}$*-complete, even when all variables are binary and the indegree of all variables is at most two.*

Result 14 ([47,45]) CondMAP-D *remains* NP*-complete on polytrees, even when all variables are binary and the indegree of all variables is at most two.*

It can be easily shown as well, using a similar argument as with the MinPE problem, that the similarly defined MinMAP-problem (given a network $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Gamma)$, an explanatory set $\mathbf{M}$, intermediate variables $\mathbf{I}$, evidence set $\mathbf{E}$ with evidence $\mathbf{e}$ and a rational number q: does $\Pr(\mathbf{m_i}, \mathbf{e}) > q$ hold *for all* joint value assignments $\mathbf{m_i}$ to $\mathbf{M}$?) is $\mathsf{co\text{-}NP}^{\mathsf{PP}}$-hard and has a $\mathsf{co\text{-}NP}^{\mathsf{PP}}$-complete decision variant; intuitively one can envisage that for solving a MinMAP-problem one needs to combine both the verification of the 'min'-property (the "co-NP-part") and probabilistic inference (the "PP-part").

Result 15 MinMAP *is* $\mathsf{co\text{-}NP}^{\mathsf{PP}}$*-hard and has a* $\mathsf{co\text{-}NP}^{\mathsf{PP}}$*-complete decision variant.*

Another problem variant, namely the *maximin a posteriori* or MmAP-problem was formulated in its decision variant as follows by De Campos and Cozman [43]: Given a probabilistic network $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Gamma)$, where $\mathbf{V}$ is partitioned into sets $\mathbf{L}$, $\mathbf{M}$, $\mathbf{I}$, and $\mathbf{E}$, and a rational number q; is there a joint value assignment $\mathbf{l}$ to $\mathbf{L}$ such that $\min_{\mathbf{m}} \Pr(\mathbf{l}, \mathbf{m}, \mathbf{e}) > q$? This problem of course resembles the Partial MAP-problem, however the set of variables is partitioned into four sets rather than three. The problem was shown $\mathsf{NP}^{\mathsf{PP}}$-hard in [43], we will show that it is in fact $\mathsf{NP}^{\mathsf{NP}^{\mathsf{PP}}}$-complete, using a reduction from the canonical $\mathsf{NP}^{\mathsf{NP}^{\mathsf{PP}}}$-complete problem EA-Majsat, defined as follows.

EA-Majsat
Instance: Let ϕ be a Boolean formula with n variables $x_i, i = 1, \ldots, n, n \geq 1$. Let $1 \leq k < l < n$, let $\mathbf{X_E}$, $\mathbf{X_A}$, and $\mathbf{X_M}$ be the sets of variables x_1 to x_k, x_{k+1} to x_l, and x_{l+1} to x_n, respectively.
Question: Is there a truth assignment to $\mathbf{X_E}$ such that for every possible truth assignment to $\mathbf{X_A}$, the majority of the truth assignments to $\mathbf{X_M}$ satisfies ϕ?

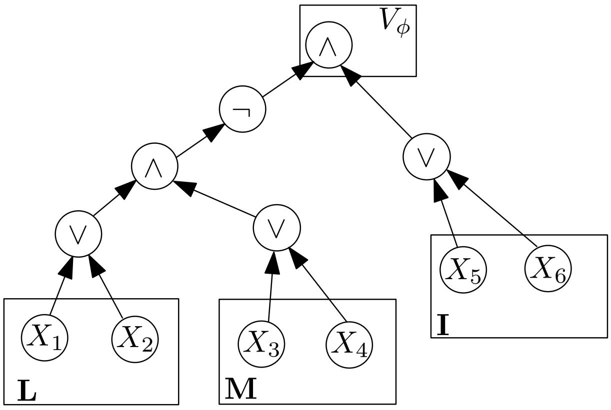

Fig. 5. The probabilistic network corresponding to $\neg((x_1 \vee x_2) \wedge (x_3 \vee x_4)) \wedge (x_5 \vee x_6)$

The intuition behind the complexity result is as follows. We have three sources of intractability which work on top of each other: marginalization over the intermediate variables (the "PP-part"), choosing $\mathbf{l}$ out of potentially exponentially many joint value assignments to $\mathbf{L}$ (one "NP-part") and verifying that $\min_{\mathbf{m}} \Pr(\mathbf{l}, \mathbf{m}, \mathbf{e}) > q$, with also potentially exponentially many joint value assignments $\mathbf{m}$ (the other "NP-part", which acts as an oracle for MinPE, but with 'yes' and 'no' answers reversed).

We construct a probabilistic network $\mathcal{B}_\phi$ from ϕ as in the hardness proof of MPE-D, however, the truth-setting part $\mathbf{X}$ is partitioned into three sets $\mathbf{L}$, $\mathbf{M}$, and $\mathbf{I}$. We take the instance $(\phi_{\text{ex}} = \neg((x_1 \vee x_2) \wedge (x_3 \vee x_4)) \wedge (x_5 \vee x_6), \mathbf{X_E} = \{x_1, x_2\}, \mathbf{X_A} = \{x_3, x_4\}, \mathbf{X_M} = \{x_5, x_6\})$ as an example; the graphical structure of the network $\mathcal{B}_{\phi_{\text{ex}}}$ constructed for ϕ_{ex} is shown in Figure 5. This EA-Majsat-instance is satisfiable: take $x_1 = x_2 =$ false, then for every truth assignment to $\{x_3, x_4\}$, the majority of the truth assignments to $\{x_5, x_6\}$ satisfy ϕ_{ex}.

Theorem 16 MmAP-D *is* $\mathsf{NP}^{\mathsf{NP}^{\mathsf{PP}}}$*-complete.*

Proof. Membership of $\mathsf{NP}^{\mathsf{NP}^{\mathsf{PP}}}$ can be proven as follows. Given a non-deterministically chosen joint value assignment $\mathbf{l}$ to $\mathbf{L}$, we can verify in polynomial time that $\min_{\mathbf{m}} \Pr(\mathbf{l}, \mathbf{m}, \mathbf{e}) > q$ using an oracle for MinMAP-D; note that $\mathsf{NP}^{\mathsf{NP}^{\mathsf{PP}}} = \mathsf{NP}^{\mathsf{co\text{-}NP}^{\mathsf{PP}}}$, as we can use an oracle for Satisfiability as an oracle for UnSatisfiability and vice versa by simply reversing the 'yes' and 'no' answers of the oracle.

To prove hardness, we show that every EA-Majsat-instance $(\phi, \mathbf{X_E}, \mathbf{X_A}, \mathbf{X_M})$ can be reduced to a corresponding instance $(\mathcal{B}_\phi, \mathbf{L}, \mathbf{M}, \mathbf{I}, \mathbf{E}, q)$ of MmAP in polynomial time. Let $\mathcal{B}_\phi$ be the probabilistic network constructed from ϕ as shown above, let $\mathbf{E} = V_\phi, \mathbf{e} =$ true and let $q = \frac{1}{2}$. Assume there exists a joint value assignment $\mathbf{l}$ to $\mathbf{L}$ such that $\min_{\mathbf{m}} \Pr(\mathbf{l}, \mathbf{m}, \mathbf{e}) > \frac{1}{2}$. Then the corresponding EA-MAJSAT-instance $(\phi, \mathbf{X_E}, \mathbf{X_A}, \mathbf{X_M})$ is satisfiable: for the truth assignment that corresponds with the joint value assignment $\mathbf{l}$, every truth assignment that corresponds to a joint value assignment $\mathbf{m}$ to $\mathbf{M}$ ensures that the majority of truth assignments to $\mathbf{X_M}$ accepts (since $\min_{\mathbf{m}} \Pr(\mathbf{l}, \mathbf{m}, \mathbf{e}) = \min_{\mathbf{m}} \sum_{\mathbf{i}} \Pr(\mathbf{l}, \mathbf{m}, \mathbf{e}, \mathbf{I} = \mathbf{i}) > \frac{1}{2}$). On the other hand, if $(\phi, \mathbf{X_E}, \mathbf{X_A}, \mathbf{X_M})$ is a satisfiable EA-MAJSAT-instance, then the proposed construction ensures that $\min_{\mathbf{m}} \Pr(\mathbf{l}, \mathbf{m}, \mathbf{e}) > \frac{1}{2}$. In other words, if we can decide arbitrary instances $(\mathcal{B}_\phi, \mathbf{L}, \mathbf{M}, \mathbf{I}, \mathbf{E}, q)$ of MMAP in polynomial time, we can decide every EA-MAJSAT-instance since the construction is obviously polynomial-time bounded, hence, MMAP-D is $\mathsf{NP}^{\mathsf{NP}^{\mathsf{PP}}}$-complete. □

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 [48,49]. 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 [50].

Kwisthout et al. [46] 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 5, however, the reduction is made from KTH-SAT (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 [51] which states that P^{PP} includes the entire Polynomial Hierarchy (PH). Here an intuitive argument for membership of the classes P^{PP} and $\mathsf{FP}^{\mathsf{PP}}$ is less easy to give; one might suggest that finding the lexicographically k-th (or, in particular, the *middle*) satisfying truth assignment to a Boolean formula is more difficult than finding the first assignment; we refer to [36] for a more exhaustive discussion.

Result 17 ([46]) KTH MPE *is* $\mathsf{FP}^{\mathsf{PP}}$*-complete and has a* P^{PP}*-complete de-*

cision variant, even if all nodes have indegree at most two.

The Kth Partial MAP-problem is even harder than that, under usual assumptions [5] in complexity theory. Kwisthout et al. proved [46] 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 $\mathbf{X_K}$ and $\mathbf{X_L}$ and an integer l, what is the lexicographically k-th satisfying truth assignment to $\mathbf{X_K}$ such that exactly l truth assignments to $\mathbf{X_L}$ satisfy ϕ?). When compared to the Kth MPE problem, we need the extra PP oracle for computing the marginal distribution of the intermediate variables, analogously to Partial MAP when compared to MPE.

Result 18 ([46]) 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.*

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 [53] showed (among other results) 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 $\mathsf{P} = \mathsf{NP}$ [53]. 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; the same holds for finding the Kth MPE.

Result 19 ([53]) MPE *cannot be approximated within any fixed ratio unless* $\mathsf{P} = \mathsf{NP}$.

Result 20 ([53]) Kth MPE *cannot be approximated within any fixed ratio unless* $\mathsf{P} = \mathsf{NP}$.

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

[5] To be more precise, the assumptions that the inclusions in the Counting Hierarchy [52] are strict.

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, as a function of the input size, or as an absolute threshold.

Result 21 (follows as a corollary from Result 2) MPE *cannot be approximated within any approximation factor* $f(n)$ *unless* $\mathsf{P} = \mathsf{NP}$.

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 [45] proved that approximating PARTIAL MAP on polytrees within a factor of $2^{|x|^\epsilon}$ is NP-hard for any fixed $\epsilon, 0 \leq \epsilon < 1$, where $|x|$ is the size of the instance.

Result 22 ([45]) PARTIAL MAP *cannot be approximated within a factor of* $2^{|x|^\epsilon}$ *for any fixed* $\epsilon, 0 \leq \epsilon < 1$*, even when restricted to polytrees, unless* $\mathsf{P} = \mathsf{NP}$.

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-q MPE, PROBABILITY-q PARTIAL MAP), treewidth (TREEWIDTH-tw MPE, TREEWIDTH-tw 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-q TREEWIDTH-tw PARTIAL MAP) and probability and number of intermediate variables (PROBABILITY-q INTERMEDIATE-l PARTIAL MAP).

Bodlaender et al. [40] 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 23 ([40]) *Any instance x of* PROBABILITY-Q MPE *can be solved in time $\mathcal{O}(f(q) \cdot |x|^c)$ for an arbitrary function f and an instance-independent constant c and is thus fixed-parameter tractable for $\{q\}$.*

Intuitively this result implies that finding the most probable explanation can be done efficiently if the probability of that explanation is high.

Sy [39] 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, and we assumed that the number of values per variable is bounded by a constant, this algorithm is exponential only in the treewidth of the network. Hence, if the treewidth tw is seen as a fixed parameter, then the algorithm runs in polynomial time.

Result 24 ([39]) *Any instance x of* TREEWIDTH-TW MPE *can be solved in time $\mathcal{O}(f(\mathrm{tw}) \cdot |x|^c)$ for an arbitrary function f and an instance-independent constant c and is thus fixed-parameter tractable for $\{\mathrm{tw}\}$.*

This result implies that finding the most probable explanation can be done efficiently also if the treewidth of the network is low.

Sy's algorithm [39] in fact finds the k most probable explanations (rather than only *the* most probable) and has a running time of $\mathcal{O}(k \cdot\ n \cdot\ \mid C \mid)$, where $\mid C \mid$ 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 25 ([39]) *Any instance x of* TREEWIDTH-TW KTH MPE *can be solved in time $\mathcal{O}(f(\mathrm{tw}, k) \cdot |x|^c)$ for an arbitrary function f and an instance-independent constant c and is thus fixed-parameter tractable for $\{tw, k\}$.*

Finding the k-th most probable explanation thus can be done efficiently if both k and the treewidth of the network are low.

In *multi-dimensional classifiers* (MBCs) one effectively solves an MPE problem in a network where the dependencies of the network are constrained: there are no arcs from the evidence set (in MBCs: feature variables) to the explanation set (in MBCs: classification variables). Observe that from the proof of Result 2 it follows that solving MBCs is NP-hard in general. Finding the most probable explanation in these restricted graphs can be done in polynomial time if both the treewidth of the evidence set and the number of classification variables are bounded (i.e., no restrictions are imposed on the topology of the explanation set) [54].

Result 26 ([54]) *Solving MBCs is fixed-parameter tractable for* $\{tw_{\mathbf{E}}, |\mathbf{M}|\}$.

Furthermore, if the MBC can be *class-bridge decomposed* [55] into components such that the maximal number of class variables $|\mathbf{C_j}|$ per component $\mathbf{C_j}$ and the treewidth of the evidence set are bounded, solving MBCs can be done in polynomial time [55].

Result 27 ([55]) *Solving MBCs is fixed-parameter tractable for* $\{tw_{\mathbf{E}}, \max_{\mathbf{C_j}} |\mathbf{C_j}|\}$.

When we consider PARTIAL MAP then restricting either the probability or the treewidth is insufficient to render the problem tractable. Park and Darwiche [45] established NP-completeness of PARTIAL MAP restricted to polytrees with at most two parents per node, i.e., networks with treewidth at most 2, yet with an unbounded number of values per variable. Recently, among other results, De Campos [56] proved NP-completeness even for binary variables, strengthening the previous result. Furthermore, it is easy to see that deciding PARTIAL MAP includes solving the INFERENCE problem, even if q, 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., [33, 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 $\Pr(C = \text{TRUE} \mid V = \text{TRUE}) = q + \epsilon$ and $\Pr(C = \text{TRUE} \mid V = \text{FALSE}) = q - \epsilon$ for an arbitrary small value ϵ. Now, $\Pr(C = \text{TRUE}) > q$ 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 q boils down to deciding INFERENCE which is PP-complete.

Result 28 ([45,56]) TREEWIDTH-TW PARTIAL MAP *is para-*NP*-complete for* $\{\text{tw}\}$.

Result 29 PROBABILITY-Q PARTIAL MAP *is para-*PP*-complete for* $\{q\}$.

However, the algorithm of Bodlaender et al. [40] can be adapted to find Partial MAPs as well [6] . The algorithm iterates over a topological sort $1, \ldots, i, \ldots, n$ of the nodes of the network. At one point, the algorithm computes $\Pr(V_{i+1} \mid \mathbf{v})$ for a particular joint value assignment $\mathbf{v}$ 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

[6] Hans L. Bodlaender, personal communication.

V_i. This can be done in polynomial time if either the number of intermediate variables (l) is fixed or the treewidth of the network (tw) is fixed.

Result 30 (adapted from [40]) *Any instance x of* PROBABILITY-Q TREEWIDTH-TW PARTIAL MAP *can be solved in time $\mathcal{O}(f(q, \text{tw}) \cdot |x|^c)$ for an arbitrary function f and an instance-independent constant c and is thus fixed-parameter tractable for $\{q, \text{tw}\}$.*

Result 31 (adapted from [40]) *Any instance x of* PROBABILITY-Q INTERMEDIATE-L PARTIAL MAP *can be solved in time $\mathcal{O}(f(q, l) \cdot |x|^c)$ for an arbitrary function f and an instance-independent constant c and is thus fixed-parameter tractable for $\{q, l\}$.*

Intuitively, 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 [21]?

We do not claim to have answers to such questions; these are to be decided by, e.g., cognitive psychologists. However, the overview of known results in this paper may aid the computational modeler in finding potential sources of intractability, i.e., parameters that render her model computationally intractable when unbounded. Whether the outcome is received as a blessing (because empirical results like reaction times and error rates may *confirm* those sources of intractability, showing that indeed the cognitive task is intractable when these parameters are unbounded, 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

The author is supported by the OCTOPUS project under the responsibility of the Embedded Systems Institute. This project is partially supported by the Netherlands Ministry of Economic Affairs under the Embedded Systems Institute program. The author wishes to thank Iris van Rooij, Hans Bodlaender, Moritz Müller, Todd Wareham and the anonymous reviewers for helpful discussions and valuable suggestions on earlier versions of this paper.

Appendix A: overview of complexity classes

In the main text of the paper, a number of complexity classes have been introduced. Here we give a short overview of these classes, composed of both a formal definition and an intuitive notion. We refer the interested reader to the classical textbooks for more background on these classes.

NP
Formal definition: Class of problems for which membership is decidable in polynomial time by a non-deterministic Turing Machine (with at least one accepting path); alternatively, class of problems for which membership is verifiable in polynomial time by a deterministic Turing Machine using a certificate.
Intuitive notion: A problem is in NP if one can easily verify membership when given a proof, e.g., one can easily verify that a Boolean formula is satisfiable, when given a satisfying truth instantiation.

co-NP
Formal definition: The complement set of NP, i.e., the class of problems for which *non*-membership is decidable in polynomial time by a non-deterministic Turing Machine; alternatively, class of problems for which membership can be falsified in polynomial time by a deterministic Turing Machine using a certificate.
Intuitive notion: A problem is in co-NP if one can easily verify non-membership when given a counterexample, e.g., one can easily falsify that a Boolean formula is a contradiction, when given a satisfying truth instantiation.

FNP

Formal definition: Class of functions that are computable in polynomial time on a non-deterministic Turing Transducer.
Intuitive notion: Strongly related to NP, but for functions rather than yes/no decision problems.

PP

Formal definition: Class of problems for which membership is decidable in polynomial time on a Probabilistic Turing Machine with an arbitrary small majority.
Intuitive notion: While for NP it suffices that there is at least *one* admissible solution, and for $\mathsf{co\text{-}NP}$ it is necessary that *no* solution [7] is admissible, for PP we demand that the (strict) *majority* of solutions are admissible. Typically associated with reasoning under uncertainty, in particular probabilistic inference.

P^{NP}, $\mathsf{FP}^{\mathsf{NP}}$

Formal definition: Class of problems for which membership is decidable on a deterministic Turing Machine, respectively functions that are computable on a deterministic Turing Transducer, in polynomial time with access to an oracle for problems in NP.
Intuitive notion: P^{NP} and $\mathsf{FP}^{\mathsf{NP}}$ are classes that are typically associated with the problem of finding an *optimal* solution, such as the lexicographically first satisfying truth instantiation.

P^{PP}, $\mathsf{FP}^{\mathsf{PP}}$

Formal definition: Class of problems for which membership is decidable on a deterministic Turing Machine, respectively functions that are computable on a deterministic Turing Transducer, in polynomial time with access to an oracle for problems in PP.
Intuitive notion: P^{PP} and $\mathsf{FP}^{\mathsf{PP}}$ are classes that are typically associated with the problem of finding the *k-th best* solution, such as the lexicographically k-th satisfying truth instantiation.

$\mathsf{NP}^{\mathsf{PP}}$, $\mathsf{co\text{-}NP}^{\mathsf{PP}}$

Formal definition: Class of problems for which membership, respectively non-membership, is decidable in polynomial time on a non-deterministic Turing Machine with access to an oracle for problems in PP.
Intuitive notion: $\mathsf{NP}^{\mathsf{PP}}$ and $\mathsf{co\text{-}NP}^{\mathsf{PP}}$ are classes that are typically associated with problems that combine probabilistic inference with either selecting candidate-solutions ($\mathsf{NP}^{\mathsf{PP}}$) or verifying properties ($\mathsf{co\text{-}NP}^{\mathsf{PP}}$).

$\mathsf{NP}^{\mathsf{NP}^{\mathsf{PP}}}$

[7] Or *every* solution, for the dual problem, compare UnSatisfiability to Tautology.

Formal definition: Class of problems for which membership is decidable in polynomial time on a non-deterministic Turing Machine with access to an oracle for problems in $\mathsf{NP}^{\mathsf{PP}}$.
Intuitive notion: The class of problems typically associated with problems that combine selecting solutions, verifying properties, and probabilistic inference, each of which individually adds to the complexity of the problem.

$\mathsf{P}^{\mathsf{NP}^{\mathsf{PP}}}$, $\mathsf{P}^{\mathsf{PP}^{\mathsf{PP}}}$, $\mathsf{FP}^{\mathsf{NP}^{\mathsf{PP}}}$, $\mathsf{FP}^{\mathsf{PP}^{\mathsf{PP}}}$
Formal definition: Class of problems for which membership is decidable, respectively functions that are computable, in polynomial time on a deterministic Turing Machine (Transducer) with access to an oracle for problems in $\mathsf{NP}^{\mathsf{PP}}$, respectively $\mathsf{PP}^{\mathsf{PP}}$.
Intuitive notion: The classes of problems that are created when augmenting P^{NP}, P^{PP}, $\mathsf{FP}^{\mathsf{NP}}$ and $\mathsf{FP}^{\mathsf{PP}}$ with an additional oracle for PP. These problems typically combine probabilistic inference with finding optimal, respectively k-th best, solutions.

FPT
Formal definition: Class of problems that have a parameter l such that membership can be decided in $\mathcal{O}(f(l) \cdot |x|^c)$ for any function f and constant c.
Intuitive notion: A problem can be hard in general, but tractable when a particular parameter of the problem instances is assumed to be fairly small. If a problem is, e.g., NP-complete for every (but finitely many) value of l, then the problem is para-NP-complete for parameter l. An example is $l-$SATISFIABILITY, where the Boolean formula is in l-CNF form. For each value of l larger than two, $l-$SATISFIABILITY is NP-complete.

Appendix B: overview of computational problems

In this appendix we formally define the relevant computational problems that are used in the paper, together with their computational complexity.

SATISFIABILITY
Instance: Let ϕ be a Boolean formula with n variables $x_i, i = 1, \ldots, n, n \geq 1$.
Question: Is there a truth assignment to $x_1, \ldots, x_n$ that satisfies ϕ?
Comment: NP-complete [57].

3SAT
Instance: As in SATISFIABILITY, but now ϕ is in 3-CNF form.
Question: Is there a truth assignment to $x_1, \ldots, x_n$ that satisfies ϕ?
Comment: NP-complete [58].

TAUTOLOGY
Instance: As in SATISFIABILITY.
Question: Does every truth assignment to $x_1, \ldots, x_n$ satisfy ϕ?
Comment: co-NP-complete (follows by definition from NP-completeness of SATISFIABILITY).

MAJSAT
Instance: As in SATISFIABILITY.
Question: Does the majority of truth assignments to $x_1, \ldots, x_n$ satisfy ϕ?
Comment: PP-complete [31].

MAJ3SAT
Instance: As in SATISFIABILITY, but now ϕ is in 3-CNF form.
Question: Does the majority of truth assignments to $x_1, \ldots, x_n$ satisfy ϕ?
Comment: PP-complete [31].

E-MAJSAT
Instance: As in SATISFIABILITY, furthermore we partition the variables into sets $\mathbf{X_E}$ and $\mathbf{X_M}$.
Question: Is there a truth assignment to $\mathbf{X_E}$ such that the majority of the truth assignments to $\mathbf{X_M}$ satisfy ϕ?
Comment: $\mathsf{NP}^{\mathsf{PP}}$-complete [59].

EA-MAJSAT
Instance: As in SATISFIABILITY, furthermore we partition the variables into sets $\mathbf{X_E}$, $\mathbf{X_A}$ and $\mathbf{X_M}$.
Question: Is there a truth assignment to $\mathbf{X_E}$ such that for every possible truth assignment to $\mathbf{X_A}$, the majority of the truth assignments to $\mathbf{X_M}$ satisfies ϕ?
Comment: $\mathsf{NP}^{\mathsf{NP}^{\mathsf{PP}}}$-complete [59].

LEXSAT
Instance: As in SATISFIABILITY.
Output: The lexicographically first truth assignment that satisfies ϕ, or $\perp$ if no such assignment exists.
Comment: $\mathsf{FP}^{\mathsf{NP}}$-complete [35].

LEXSAT-D
Instance: As in SATISFIABILITY.
Question: Is the least significant bit in the lexicographically first truth assignment that satisfies ϕ odd?
Comment: P^{NP}-complete [35].

KTHSAT
Instance: As in SATISFIABILITY, natural number k.
Output: The lexicographically k-th truth assignment that satisfies ϕ, or $\perp$

if no such assignment exists.
Comment: $\mathsf{FP}^{\mathsf{PP}}$-complete [36].

KthSat-D
Instance: As in Satisfiability, natural number k.
Question: Is the least significant bit in the lexicographically k-th truth assignment that satisfies ϕ odd?
Comment: P^{PP}-complete [36].

KthNumSat
Instance: As in E-Majsat, natural numbers k and l.
Output: The lexicographically k-th truth assignment to the set $\mathbf{X_E}$ for which exactly l truth assignments to $\mathbf{X_M}$ satisfy ϕ, or $\perp$ if no such assignment exists.
Comment: $\mathsf{FP}^{\mathsf{PP}^{\mathsf{PP}}}$-complete [33].

MPE
Instance: A probabilistic network $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Gamma)$, 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 $\mathbf{M}$ and evidence $\mathbf{e}$, or $\perp$ if $\Pr(\mathbf{M} = \mathbf{m}, \mathbf{E} = \mathbf{e}) = 0$ for every joint value assignment $\mathbf{m}$ to $\mathbf{M}$.
Comment: $\mathsf{FP}^{\mathsf{NP}}$-complete [46].

MPE-D
Instance: As in MPE, rational number q.
Question: Is there a joint value assignment $\mathbf{m}$ to the nodes in $\mathbf{M}$ with evidence $\mathbf{e}$ with probability $\Pr(\mathbf{M} = \mathbf{m}, \mathbf{E} = \mathbf{e}) > q$?
Comment: NP-complete [44].

MPEe-D
Instance: As in MPE, rational number q.
Question: Is there a joint value assignment $\mathbf{m}$ to the nodes in $\mathbf{M}$ with evidence $\mathbf{e}$ with probability $\Pr(\mathbf{M} = \mathbf{m} \mid \mathbf{E} = \mathbf{e}) > q$?
Comment: PP-complete [43].

MPE-D′
Instance: As in MPE, designated variable $M \in \mathbf{M}$ with designated value m.
Question: Does M have the value m in the most probable joint value assignment $\mathbf{m}$ to $\mathbf{M}$?
Comment: P^{NP}-complete [33].

is-an-MPE
Instance: As in MPE, joint value assignment $\mathbf{m}$ to $\mathbf{M}$.
Question: Is $\mathbf{m}$ the most probable joint value assignment to $\mathbf{M}$?

Comment: co-NP-complete [40].

better-MPE
Instance: As in MPE, joint value assignment $\mathbf{m}$ to $\mathbf{M}$.
Output: A joint value assignment $\mathbf{m'}$ to $\mathbf{M}$ such that $\Pr(\mathbf{M} = \mathbf{m'} \mid \mathbf{E} = \mathbf{e}) > \Pr(\mathbf{M} = \mathbf{m} \mid \mathbf{E} = \mathbf{e})$, or $\perp$ if no such joint value assignment exists.
Comment: NP-hard [40].

MinPE-D
Instance: As in MPE, rational number q.
Question: Does $\Pr(\mathbf{M} = \mathbf{m_i}, \mathbf{E} = \mathbf{e}) > q$ hold *for all* joint value assignments $\mathbf{m_i}$ to $\mathbf{M}$?
Comment: co-NP-complete (Section 4).

Kth MPE
Instance: As in MPE, natural number k.
Output: The k-th most probable joint value assignment $\mathbf{m}$ to the nodes in $\mathbf{M}$ and evidence $\mathbf{e}$, or $\perp$ if k is larger than the number of joint value assignments $\mathbf{m}$ to $\mathbf{M}$ for which $\Pr(\mathbf{m}, \mathbf{e}) > 0$.
Comment: $\mathsf{FP}^{\mathsf{NP}}$-complete [46].

Kth MPE-D
Instance: As in MPE, natural number k, designated variable $M \in \mathbf{M}$ with designated value m.
Question: Does M have the value m in the k-th most probable joint value assignment $\mathbf{m}$ to the nodes in $\mathbf{M}$ and evidence $\mathbf{e}$?
Comment: P^{NP}-complete [46].

Partial MAP
Instance: A probabilistic network $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Gamma)$, 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 $\mathbf{M}$ and evidence $\mathbf{e}$, or $\perp$ if $\Pr(\mathbf{m}, \mathbf{e}) = 0$ for every joint value assignment $\mathbf{m}$ to $\mathbf{M}$.
Comment: $\mathsf{FP}^{\mathsf{NP}^{\mathsf{PP}}}$-complete [46].

Partial MAP-D
Instance: As in Partial MAP, rational number q.
Question: Is there a joint value assignment $\mathbf{m}$ to the nodes in $\mathbf{M}$ with evidence $\mathbf{e}$ with probability $\Pr(\mathbf{M} = \mathbf{m}, \mathbf{E} = \mathbf{e}) > q$?
Comment: $\mathsf{NP}^{\mathsf{PP}}$-complete [45].

CondMAP-D
Instance: As in Partial MAP, rational number q.

Question: Is there a joint value assignment $\mathbf{m}$ to $\mathbf{M}$ such that $\Pr(\mathbf{E} = \mathbf{e} \mid \mathbf{M} = \mathbf{m}) > q$?
Comment: $\mathsf{NP}^{\mathsf{PP}}$-complete [45,47].

MinMAP-D
Instance: As in Partial MAP, rational number q.
Question: Is $\Pr(\mathbf{M} = \mathbf{m_i}, \mathbf{E} = \mathbf{e}) > q$ for all joint value assignments $\mathbf{m_i}$ to $\mathbf{M}$?
Comment: $\mathsf{co\text{-}NP}^{\mathsf{PP}}$-complete (Section 5).

MmAP-D
Instance: A probabilistic network $\mathcal{B} = (\mathbf{G}_{\mathcal{B}}, \Gamma)$, 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 explanation sets $\mathbf{L}$ and $\mathbf{M}$, rational number q.
Question: Is there a joint value assignment $\mathbf{l}$ to $\mathbf{L}$ such that $\min_{\mathbf{m}} \Pr(\mathbf{L} = \mathbf{l}, \mathbf{M} = \mathbf{m}, \mathbf{E} = \mathbf{e}) > q$?
Comment: $\mathsf{NP}^{\mathsf{NP}^{\mathsf{PP}}}$-complete (Section 5).

Kth Partial MAP
Instance: As in Partial MAP, natural number k.
Output: The k-th most probable joint value assignment $\mathbf{m}$ to the nodes in $\mathbf{M}$ and evidence $\mathbf{e}$, or $\perp$ if k is larger than the number of joint value assignments $\mathbf{m}$ to $\mathbf{M}$ for which $\Pr(\mathbf{m}, \mathbf{e}) > 0$.
Comment: $\mathsf{FP}^{\mathsf{PP}^{\mathsf{PP}}}$-complete [46].

Kth Partial MAP-D
Instance: As in Partial MAP, designated variable $M \in \mathbf{M}$ with designated value m.
Question: Does M have the value m in the k-th most probable joint value assignment $\mathbf{m}$ to the nodes in $\mathbf{M}$ and evidence $\mathbf{e}$?
Comment: $\mathsf{P}^{\mathsf{PP}^{\mathsf{PP}}}$-complete [46].

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) Bayesian Networks: a decision tool to improve portfolio risk analysis. *Journal of applied finance*, **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] Sticha, P. J., Buede, D. M., and Rees, R. L. (2006) Bayesian model of the effect of personality in predicting decisionmaker behavior. In Van der Gaag, L. C. and Almond, R. (eds.), *Proceedings of the Fourth Bayesian Modelling Applications Workshop*.

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

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

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

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

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

[17] 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.

[18] 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.

[19] Kwisthout, J., Wareham, T., and van Rooij, I. (in press) Letter to the Editor: Bayesian Intractability is Not an Ailment Approximation Can Cure. *Cognitive Science.*

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

[21] Blokpoel, M., Kwisthout, J., van der Weide, T. P., and van Rooij, I. (2010) How action understanding can be rational, Bayesian *and* tractable. In Ohlsson, S. and Catrambone, R. (eds.), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*, pp. 1643–1648. Cognitive Science Society, Austin, TX.

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

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

[24] Koller, D. and Friedman, N. (2009) *Probabilistic Graphical Models: principles and techniques.* MIT press, Cambridge, MA.

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

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

[27] 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.

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

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

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

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

[32] Littman, M. L., Majercik, S. M., and Pitassi, T. (2001) Stochastic boolean satisfiability. *Journal of Automated Reasoning*, **27**, 251–296.

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

[34] Littman, M. L., Goldsmith, J., and Mundhenk, M. (1998) The computational complexity of probabilistic planning. *Journal of Artificial Intelligence Research*, **9**, 1–36.

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

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

[37] Gasarch, W.I., Krentel, M.W., and Rappoport, K.J. (1995) OptP as the normal behavior of NP-Complete problems. *Mathematical Systems Theory*, **28(6)**, 487-514.

[38] Flum, J. and Grohe, M. (2006) *Parameterized Complexity Theory.* Springer Verlag, Berlin.

[39] 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.

[40] 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.

[41] 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.

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

[43] De Campos, C. P. and Cozman, F. G. (2005) The inferential complexity of Bayesian and credal networks. *Proceedings of the Nineteenth international joint conference on Artificial intelligence*, pp. 1313–1318.

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

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

[46] Kwisthout, J., Bodlaender, H. L. and Van der Gaag, L. C. (2011) The complexity of finding kth most probable explanations in probabilistic networks. In Cerná, I., Gyimóthy, T., Hromkovic, J., Jefferey, K., Královic, R., Vukolic, M., and Wolf, S. (eds.), *Proceedings of the 37th International Conference on Current Trends in Theory and Practice of Computer Science*, LNCS 6543, pp. 356–367. Springer, Berlin.

[47] Van der Gaag, L. C., Bodlaender, H. L., and Feelders, A. J. (2004) Monotonicity in Bayesian networks. In Chickering, M. and Halpern, J. (eds.), *Proceedings of the Twentieth Conference on Uncertainty in Artificial Intelligence*, pp. 569–576. Arlington: AUAI press.

[48] 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.

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

[50] 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.

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

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

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

[54] De Waal, P. R. and Van der Gaag, L. C. (2007) Inference and Learning in multi-dimensional Bayesian network classifiers. In Mellouli, K. (ed.) *Proceedings of the Ninth European Conference on Symbolic and Quantitative Approaches to Reasoning under Uncertainty*, pp. 501–511.

[55] Bielza, C., Li, G., and Larrañga, P. (2011) Multi-dimensional classification with Bayesian networks. *International Journal of Approximate Reasoning*, **52**, 705-727.

[56] De Campos, C. P. (2011) New Complexity Results for MAP in Bayesian Networks. In Walsh, T. (ed.) *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence*, pp. 2100–2106.

[57] Cook, S. A. (1971) The complexity of theorem proving procedures. *Proceedings of the Third Annual ACM Symposium on Theory of Computing*, pp. 151–158.

[58] Karp, R. M. (1972) Reducibility among combinatorial problems. In Miller, R.E. and Thatcher, J.W. (eds.) *Complexity of Computer Computations*, pp. 85-103. Plenum, New York, NY.

[59] Wagner, K. W. (1986) The complexity of combinatorial problems with succinct input representation. *Acta Informatica*, **23**, 325–356.