

Two new notions of abduction in Bayesian networks

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

November 17, 2010

Abstract

Most Probable Explanation and (Partial) MAP are well-known problems in Bayesian networks that correspond to Bayesian or probabilistic inference of the most probable explanation of observed phenomena, given full or partial evidence. These problems have been studied extensively, both from a knowledge-engineering starting point (see [15] for an overview) as well as a complexity-theoretic point of view (see [13] for an overview). Algorithms, both exact and approximate, are studied in e.g. [21, 25, 19, 28]. In this paper, we introduce two new notions of abduction-like problems in Bayesian networks, motivated from cognitive science, namely the problem of finding the most *simple* and the most *informative* explanation for a set of variables, given evidence. We define and motivate these problems, show that these problems are computationally intractable in general, but become tractable when some particular constraints are met.

1 Introduction

Abduction (the process of finding a suitable explanation of observed phenomena) is, according to the well-known philosopher Charles Sanders Pierce¹ perhaps the most essential mechanism through which we acquire knowledge or information. While since long obvious in fields like medicine (“which disease is causing these symptoms”), history (“which circumstances led to this uprising”), and the sciences (“what theory explains this behavior”), the problem becomes also more and more important in fields as cognitive science, e.g., in computational models of goal inference [1], action understanding [4], or vision [29].

In Bayesian networks, the abduction problem is often encoded as the problem of determining the joint value assignment to a set of variables in the

¹See [6] for an overview of Pierce’s work on abduction.

network which has the highest posterior probability given the observed values of other variables in the network. This problem is known as MPE or (PARTIAL) MAP in the literature². The abduction problem in Bayesian networks (be it formalized as MPE or PARTIAL MAP) is intractable in general, yet can be constrained to make computation tractable [13]. For example, the MPE problem, where the network is bi-partitioned into a set of variables for which the most probable joint value assignment is sought (the *explanation set*) and a set of variables whose values are observed (the *evidence*), can be solved fast if the most probable explanation is very likely (and thus, correspondingly, competing explanations are very improbable) [3]. However, the bi-partition constraint is often not satisfied in practice or may lead to unexpected results, forcing us to solve a PARTIAL MAP problem (where only partial evidence is available, i.e., not all variables outside the explanation set are observed). For example, assume our domain consists of four variables: the classification variable *Illness* (I) with values *pneumonia* and *influenza*, and the three variables *Headaches* (H), *Fever* (F) and *Shortness-of-breath* (S), all with values TRUE and FALSE. Assume the following conditional probabilities are known:

$$\begin{aligned} \Pr(I = \text{pneumonia}, H = \text{TRUE} | S = \text{TRUE}, F = \text{TRUE}) &= 0.25 \\ \Pr(I = \text{pneumonia}, H = \text{FALSE} | S = \text{TRUE}, F = \text{TRUE}) &= 0.30 \\ \Pr(I = \text{influenza}, H = \text{TRUE} | S = \text{TRUE}, F = \text{TRUE}) &= 0.35 \\ \Pr(I = \text{influenza}, H = \text{FALSE} | S = \text{TRUE}, F = \text{TRUE}) &= 0.10 \end{aligned}$$

Furthermore assume that we observe $S = \text{TRUE}$ and $F = \text{TRUE}$, but we do not know whether the patient has headaches or not. However, we are interested in the most probable value of I , rather than the most probable joint value assignment to $\{I, H\}$. From the conditional probabilities we observe that in the latter case $I = \text{influenza}, H = \text{TRUE}$ is the most probable explanation of $S = \text{TRUE}, F = \text{TRUE}$. When we look only at I and marginalize over H , we compute $\Pr(I = \text{pneumonia} | S = \text{TRUE}, F = \text{TRUE}) = 0.55$ and $\Pr(I = \text{influenza} | S = \text{TRUE}, F = \text{TRUE}) = 0.45$, in other words, the most probable illness given the observed signs is pneumonia.

The merit of giving results that intuitively better match our notion of abduction (i.e., solving a PARTIAL MAP problem rather than a MPE problem) comes at the price of higher computational demands [23, 20, 13]. Moreover, when used as a computational framework of cognitive tasks in which abduction plays a central role, it is questionable whether PARTIAL MAP is actually the most plausible underlying computational problem in such cognitive tasks. We

²Note that other definitions of “inference to the best explanation” (IBE) other than the “most probable explanation” have been advocated, for example picking the explanation \mathbf{m} which has the maximum likelihood $\Pr(\mathbf{E} = \mathbf{e} | \mathbf{M} = \mathbf{m})$ of all possible joint value assignments to \mathbf{M} [11], using parsimonious covering theory to select the least irredundant explanation [22], or using a coherence measure $C(\mathbf{M}, \mathbf{E}) = \frac{\Pr(\mathbf{M} \wedge \mathbf{E})}{\Pr(\mathbf{M} \vee \mathbf{E})}$ [9] to rank explanations. Here, however, our starting point will be the “most probable explanation” stance, while this appears to be the dominant way to define Bayesian abduction within the Bayesian network community; see [10] for a discussion of alternative notions of IBE.

will give two examples.

Example 1.1. *Mr. Jones typically comes to work by train. Today Mr. Jones is late while he has been seen to leave his house at the usual time. One explanation can be that the train is delayed. However, it might also be the case that Mr. Jones was the unlucky individual who walked through 11th Street at 8.03 AM and was shot during an armed bank robbery. When trying to explain why Mr. Jones is not at his desk on 8.30 AM, there are a number of variables we might take into account, for example whether he has to change trains. A whole lot of variables are typically not taken into account because they are normally not relevant in most of the cases, for example whether Mr. Jones walked on the left or right pavement in 11th Street. Only in the awkward coincidence that Mr. Jones was in the wrong place at the wrong time they become relevant to explain why he is not at work.*

Is it plausible that, in order to determine the most likely cause of Mr. Jones' lateness, we really marginalize over all these unknown—and usually irrelevant—variables in order to infer the most likely explanation? Or do we actually think a train delay is a likely explanation because it explains his absence for the overwhelming majority of values for these variables, save that particular situation where he was walking on the left pavement and passed the bank at 8.03 AM while there was a shooting at that very moment? We will argue that the latter may be a more plausible cognitive model, congruent with recent findings in cognitive science.

An other potential problem when using PARTIAL MAP as the underlying computational model is that the choice of the explanation set is crucial: we seek to find the most probable joint value assignment to that set of variables, i.e., the explanation set is given in the input. The question *what constitutes a good explanation* will be illustrated in the following example.

Example 1.2. *Dr. Brown is examining Mr. Smith. She notices an increased body temperature and Mr. Smith complains of shortness of breath, coughing with phlegm, and pain while breathing. A correct, but hardly informative, explanation of these signs is 'Mr. Smith is ill.' This explanation has a higher probability (say 0.998) than the much more informative explanation 'Mr. Smith has pneumonia.' which may have a probability of 0.96. The latter explanation of course has more explanatory power at the cost of little probability mass, and is thus preferred over the former although this explanation has a higher probability.*

Note that there are many situation-specific circumstances that may determine whether a more specific explanation is needed. While a general practitioner will need an explanation that is specific enough to successfully describe medication, Mr. Smith's project manager needs only a general explanation why he won't be at his desk for some time. Sometimes it might be costly to determine more specific explanations. Typically there is a trade-off between specificity and probability and the impact of making the *wrong* decision is crucial in determining the probability threshold.

In this paper we introduce two new notions of abduction in Bayesian networks, namely finding the explanation that *assumes as little as possible of the values of other variables that are not observed* (which we will denote as the MOST SIMPLE EXPLANATION problem), and finding the explanation that *carries the most information given some probability threshold* (which we will denote as the MOST INFORMATIVE EXPLANATION problem). After the introduction of some preliminaries in Section 2, the problem variants will be formalized in Section 3 and 4, respectively. We will give complexity results for the problems in general and show under which constraints these problems are feasible. We conclude the paper in Section 5.

2 Preliminaries

A Bayesian or probabilistic network \mathcal{B} is a graphical structure that models a set of stochastic variables, the (in-)dependencies among these variables, and a joint probability distribution over these variables. \mathcal{B} includes a directed acyclic graph $\mathbf{G} = (\mathbf{V}, \mathbf{A})$, modeling the variables and (in-) dependencies in the network, and a set of parameter probabilities Γ in the form of conditional probability tables (CPTs), capturing the strengths of the relationships between the variables. The network models a joint probability distribution $\Pr(\mathbf{V}) = \prod_{i=1}^n \Pr(V_i | \pi(V_i))$ over its variables, where $\pi(V_i)$ denotes the parents of V_i in \mathbf{G} . We will use upper case letters to denote individual nodes in the network, upper case bold letters to denote sets of nodes, lower case letters to denote value assignments to nodes, and lower case bold letters to denote joint value assignments to sets of nodes. Every (posterior) probability of interest in Bayesian networks can be computed using well known lemmas in probability theory, like Bayes' theorem ($\Pr(H | E) = \frac{\Pr(E|H)\Pr(H)}{\Pr(E)}$), marginalization ($\Pr(H) = \sum_{g_i} \Pr(H \wedge G = g_i)$), and the factorization property of Bayesian networks ($\Pr(\mathbf{V}) = \prod_{i=1}^n \Pr(V_i | \pi(V_i))$). More background can be found in textbooks like [21] and [12].

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 [7]. In addition to these basic concepts, to describe the complexity of various problems we will use the *probabilistic* classes PP and BPP, oracles, and fixed-parameter tractability.

The class PP contains languages L accepted in polynomial time by a *Probabilistic Turing Machine*. Such a machine augments the more traditional non-deterministic Turing Machine with a probability distribution associated with each state transition, e.g., by providing the machine with a tape, randomly filled with symbols [8]. If all choice points are binary and the probability of each transition is $\frac{1}{2}$, then the *majority* of the computation paths accept a string s if and only if $s \in L$. This majority, however, is not fixed and may (exponentially) depend on the input, e.g., a problem in PP may accept 'yes'-instances with size n with probability $\frac{1}{2} + \frac{1}{2^n}$. This makes problems in PP intractable in general, in contrast to the related complexity class BPP which is associated

with problems which allow for efficient randomized computation³. BPP accepts ‘yes’-inputs with a *bounded*⁴ majority (say $\frac{3}{4}$). This means we can amplify the probability of a correct answer arbitrary close to one by running the algorithm a polynomial amount of times and taking a majority vote on the outcome. This approach fails for unbounded majorities as $\frac{1}{2} + \frac{1}{2^n}$ as allowed by the class PP: here an exponential number of simulations (with respect to the input size) is needed to meet a constant threshold on the probability of answering correctly.

The canonical PP-complete problem is MAJSAT: given a Boolean formula ϕ , does the majority of the truth instantiations satisfy ϕ ? Indeed it is easily shown that MAJSAT encodes 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 satisfy ψ if and only if ϕ is satisfiable, thus $\text{NP} \subseteq \text{PP}$. In probabilistic networks, the INFERENCE problem of determining whether the probability $\Pr(\mathbf{X} = \mathbf{x}) > q$ is PP-complete.

A Turing Machine \mathcal{M} has *oracle access* to languages in the class A, denoted as \mathcal{M}^A , if it can “query the oracle” in one state transition, i.e., in $\mathcal{O}(1)$. We can regard the oracle as a ‘black box’ that can answer membership queries in constant time. For example, NP^{PP} is defined as the class of languages which are decidable in polynomial time on a non-deterministic Turing Machine with access to an oracle deciding problems in PP. Informally, computational problems related to probabilistic networks that are in NP^{PP} typically combine some sort of *selecting* with *probabilistic inference*. The canonical NP^{PP} -complete satisfiability variant is E-MAJSAT: given a formula ϕ whose variables are partitioned into sets $x_1 \dots x_k$ and $x_{k+1} \dots x_n$, is there an instantiation to $x_1 \dots x_k$ such that the majority of the instantiations to $x_{k+1} \dots x_n$ satisfy ϕ ?

A problem is called *fixed parameter tractable* [5] for a parameter k if it can be solved in time, exponential *only* in k and polynomial in the input size n , i.e., when the running time is $\mathcal{O}(f(k) \cdot n^c)$ for an arbitrary function f and a constant c , independent of n . In practice, this means that problem instances can be solved efficiently, even when the problem is NP-hard in general, if k is known to be small. If an NP-hard problem Π is *fixed parameter tractable* for a parameter k then k is denoted a *source of complexity* [26] of Π : bounding k renders the problem tractable, whereas leaving k unbounded ensures intractability under usual complexity-theoretic assumptions like $\text{P} \neq \text{NP}$. Such a source of complexity in computational problems that stem from Bayesian networks is often the *treewidth*[2] of the network under consideration.

To conclude this preliminaries section, we define decision variants of the MPE and PARTIAL MAP problems as follows⁵:

MPE

Instance: A probabilistic network $\mathcal{B} = (\mathbf{G}, \Gamma)$, where \mathbf{V} is partitioned into a

³Indeed BPP is assumed to be equal to P by many complexity theorists.

⁴To be precise: polynomially bounded in the input size.

⁵Observe that we consistently use marginal probabilities $\Pr(\mathbf{m}, \mathbf{e})$ here rather than conditional probabilities $\Pr(\mathbf{m} | \mathbf{e})$. This is due to the different complexity results for MPE defined with marginal and conditional probabilities [13].

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 $\Pr(\mathbf{m}, \mathbf{e}) > q$?

PARTIAL MAP

Instance: A probabilistic network $\mathcal{B} = (\mathbf{G}, \Gamma)$, where \mathbf{V} is partitioned into a set of evidence nodes \mathbf{E} with a joint value assignment \mathbf{e} , an explanation set \mathbf{M} , and a set of intermediate variables \mathbf{I} ; 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 $\Pr(\mathbf{m}, \mathbf{e}) > q$?

MPE has been proven NP-complete by Shimony [23] and is fixed parameter tractable in graphs with bounded treewidth and when the most probable explanation has a high probability [25, 3]. PARTIAL MAP has been proven to be NP^{PP} -complete by Park and Darwiche [20]. In contrast with MPE, PARTIAL MAP remains intractable when the network has bounded treewidth or when the most probable explanation has a high probability, yet is fixed parameter tractable when *both* conditions are met [13].

3 Most Simple Explanations

In this section we will introduce Most Simple Explanations as an alternative notion of Bayesian abduction with partial evidence. Our approach is partially inspired by the *Decision by Sampling* model of Stewart et al. [24], where the authors “...do not assume that people have stable, long-term internal scales along which they represent value, probability, temporal duration, or any other magnitudes.”, but, “...assume that people can only sample items from memory and then judge whether a target value is larger or smaller than these items.” [24, p. 2]. This was recently confirmed by Vul et al. [27]: few samples from a probability distribution are often adequate to make reasonably well decisions.

Based on the intuitive observation that we often do not take information into account (for example, whether Mr. Jones was walking on the left or right pavement) unless we have reason to believe that this might influence the most probable explanation, we suggest that the PARTIAL MAP problem, in which marginalization of the intermediate variables takes place, might not be the formalism that describes human behavior best in such circumstances. Instead, we suggest that one does not choose the *most probable* explanation—based on a marginalization of potentially many intermediate variables—but one rather picks the *most simple* explanation, i.e., the explanation that assumes as little as possible of the values of the intermediate variables. This can be seen as an implementation of *Occam’s razor* within the context of Bayesian abduction. We formalize this as follows: select the joint value assignment to the explanation set that is the *most probable explanation* for the *majority* of joint value assignments to the intermediate variables.

To stick with the *Mr. Jones*-example: for almost every joint value assignment to the intermediate variables Time-Of-Passing, Shooting-at-Bank, Pavement-walking-on, Bus-17-riding-by, Color-of-Mr.Jones'-suit etcetera, the explanation “Train-is-delayed = TRUE” would be the most probable explanation of Mr. Jones’ absence at work. That would be the most *simple* explanation assuming in general no particular value assignments to the color of Mr. Jones’ suit or whether bus 17 or 29 was riding by. Only for a particular combination of circumstances, the most probable explanation would shift to “Mr.Jones-got-shot-on his-way-to-work = TRUE”. Thus, when picking a few arbitrary *samples* out of this vast domain space, chances are high that all of them have “Train-is-delayed = TRUE” as the most probable explanation.

Typically, such a decision might be difficult to make when there are multiple competing alternatives. In such cases, apart from solving MPE problems, we may have a difficult task in deciding upon the majority. This is reflected in the computational complexity result of MOST SIMPLE EXPLANATION.

3.1 Computational Complexity

We will prove that MOST SIMPLE EXPLANATION is NP^{PP} -complete, and thus in the same complexity class as PARTIAL MAP. First we give a decision variant of the MOST SIMPLE EXPLANATION problem.

MOST SIMPLE EXPLANATION

Instance: A probabilistic network $\mathcal{B} = (\mathbf{G}, \Gamma)$, where \mathbf{V} is partitioned into a set of evidence nodes \mathbf{E} with a joint value assignment \mathbf{e} , an explanation set \mathbf{M} , and intermediate variables \mathbf{I} ; a rational number $0 \leq q < 1$.

Question: Is there a joint value assignment \mathbf{m} to the nodes in \mathbf{M} such that for the majority of joint value assignments \mathbf{i} to \mathbf{I} , $\Pr(\mathbf{m}, \mathbf{i}, \mathbf{e}) > q$?

With respect to the *Mr. Jones* example, the evidence \mathbf{e} would contain observations as “Has-left-house-at-usual-time = TRUE” and “Is-absent-from-work = TRUE”, the explanation set \mathbf{M} would contain, among others, the variables Train-is-delayed and Mr.Jones-got-shot-on his-way-to-work, and the set of intermediate variables would contain (among *lots of* other variables) Time-Of-Passing, Shooting-at-Bank, Pavement-walking-on, Bus-17-riding-by, Color-of-Mr.Jones'-suit etcetera. The MOST SIMPLE EXPLANATION problem, in its decision variant, asks whether there is an explanation whose posterior probability exceeds a threshold q for the majority of joint value assignments to the variables Time-Of-Passing etcetera.

To prove NP^{PP} -completeness of this problem, we construct a probabilistic network \mathcal{B}_ϕ from an E-MAJSAT instance $(\phi, \mathbf{E}, \mathbf{M})$, where ϕ is a Boolean formula with n variables, partitioned into sets $\mathbf{E} = x_1 \dots x_l$ and $\mathbf{M} = x_{l+1} \dots x_n$ for some number $0 \leq l \leq n$. For each propositional variable x_i in ϕ , a binary stochastic variable X_i is added to \mathcal{B}_ϕ , with possible values TRUE and FALSE and a uniform probability distribution. These stochastic variables in \mathcal{B}_ϕ are bi-partitioned into sets $\mathbf{X}_\mathbf{E}$ and $\mathbf{X}_\mathbf{M}$ according to the partition of ϕ . For each

logical operator in ϕ , an additional binary variable in \mathcal{B}_ϕ is introduced, whose parents are the variables that correspond to the input of the operator, and whose conditional probability table is equal to the truth table of that operator. For example, the value TRUE of a stochastic variable mimicking an *and*-operator would have a conditional probability of 1 if and only if both its parents have the value TRUE, and 0 otherwise. The variable associated with the top-level operator in ϕ is denoted as V_ϕ . Figure 1 shows the graphical structure of the probabilistic network constructed for the E-MAJSAT instance $(\phi_{\text{ex}}, \mathbf{E}, \mathbf{M})$, where $\phi_{\text{ex}} = \neg(x_1 \vee x_2) \wedge (x_3 \vee x_4)$, $\mathbf{E} = \{x_1, x_2\}$, and $\mathbf{M} = \{x_3, x_4\}$.




Figure 1: Example of construction of a probabilistic network $\mathcal{B}_{\phi_{\text{ex}}}$ for the Boolean formula $\phi_{\text{ex}} = \neg(x_1 \vee x_2) \wedge (x_3 \vee x_4)$

Theorem 3.1. MOST SIMPLE EXPLANATION is NPP^{PP} -complete.

Proof. Membership in NPP^{PP} follows from the following algorithm: non-deterministically guess a value assignment \mathbf{m} , count the joint value assignments \mathbf{i} to \mathbf{I} such that $\Pr(\mathbf{m}, \mathbf{i}, \mathbf{e}) > q$ and decide whether this is a majority. This can be done using a non-deterministic algorithm with an oracle for problems in PP, since deciding whether there *exists* a joint value assignment \mathbf{i} such that $\Pr(\mathbf{m}, \mathbf{i}, \mathbf{e}) > q$ is in NP, and thus deciding whether this is the case for the *majority* of joint value assignments is thus in PP by definition⁶.

To prove NPP^{PP} -hardness, we reduce MOST SIMPLE EXPLANATION from E-MAJSAT. We fix $q = \frac{1}{2}$. Let $(\phi, \mathbf{E}, \mathbf{M})$ be an instance of E-MAJSAT and let \mathcal{B}_ϕ be the network constructed from that instance as shown above. If there exists a satisfying solution to $(\phi, \mathbf{E}, \mathbf{M})$, then there is a truth instantiation to \mathbf{E} such that for the majority of truth instantiations to \mathbf{M} , ϕ is satisfied. For any particular corresponding joint value assignment $\mathbf{x}_\mathbf{E}$ to $\mathbf{X}_\mathbf{E}$ in \mathcal{B}_ϕ , $\Pr(V_\phi = \text{TRUE}, \mathbf{x}_\mathbf{E}, \mathbf{x}_\mathbf{M}) = 1$ for a particular joint value assignment $\mathbf{x}_\mathbf{M}$ to $\mathbf{X}_\mathbf{M}$, if and

⁶Recall that a problem is in NP if there is a Turing machine which, when presented a positive instance, has at least one accepting computation path, and in PP if the majority of computation paths accept.

only if the corresponding truth instantiation to \mathbf{M} satisfies ϕ , and 0 otherwise. Correspondingly, if $(\phi, \mathbf{E}, \mathbf{M})$ is satisfiable, then $\Pr(V_\phi = \text{TRUE}, \mathbf{x}_\mathbf{E}, \mathbf{x}_\mathbf{M}) = 1$ for the majority of joint value assignments $\mathbf{x}_\mathbf{M}$ to $\mathbf{X}_\mathbf{M}$, and thus the MOST SIMPLE EXPLANATION instance $(\mathcal{B}_\phi, V_\phi = \text{TRUE}, \mathbf{X}_\mathbf{E})$ is satisfied.

Now assume $(\mathcal{B}_\phi, V_\phi = \text{TRUE}, \mathbf{X}_\mathbf{E})$ is satisfied, i.e., there exists a joint value assignment $\mathbf{x}_\mathbf{E}$ to $\mathbf{X}_\mathbf{E}$ such that for the majority of joint value assignments $\mathbf{x}_\mathbf{M}$ to $\mathbf{X}_\mathbf{M}$, $\Pr(V_\phi = \text{TRUE}, \mathbf{x}_\mathbf{E}, \mathbf{x}_\mathbf{M}) > \frac{1}{2}$. Then there exists a truth instantiation to \mathbf{E} (namely the truth instantiation that corresponds to the joint value assignment $\mathbf{x}_\mathbf{E}$) such that the majority of truth instantiations to \mathbf{M} satisfies ϕ , and thus the E-MAJSAT instance $(\phi, \mathbf{E}, \mathbf{M})$ is satisfied. Since the reduction can obviously be done in polynomial time, this proves that MOST SIMPLE EXPLANATION is NP^{PP} -complete. \square

Note that MOST SIMPLE EXPLANATION remains NP^{PP} -complete, even if all variables are binary, each node has at most two incoming arcs, and the probability distribution of each variable is either uniform or deterministic. If $\mathbf{I} = \emptyset$ then MOST SIMPLE EXPLANATION degenerates to MPE and hence is NP -complete. If $\mathbf{M} = \emptyset$ then MOST SIMPLE EXPLANATION remains PP -complete via a reduction of MAJSAT, using a similar construction as above. With $\mathbf{M} = \emptyset$, the MOST SIMPLE EXPLANATION problem degenerates to the question whether for the majority of joint value assignments \mathbf{i} to \mathbf{I} , $\Pr(\mathbf{i}, \mathbf{e}) > q$; this problem is in PP by a similar argument as above. We set again the evidence to be $V_\phi = \text{TRUE}$ and fix $q = \frac{1}{2}$. Now, $\Pr(\mathbf{i}, \mathbf{e}) = 1$ if \mathbf{i} corresponds to a satisfying truth instantiation to ϕ , and 0 otherwise. Thus, if $\Pr(\mathbf{i}, \mathbf{e}) > \frac{1}{2}$ for the *majority* of joint value assignments \mathbf{i} , then the majority of truth instantiations satisfy ϕ and hence MOST SIMPLE EXPLANATION is PP -complete.

3.2 A Randomized Algorithm

A deterministic algorithm for MOST SIMPLE EXPLANATION might iterate over \mathbf{m} and count the joint value assignments \mathbf{i} to \mathbf{I} for which $\text{argmax}_\mathbf{M} \Pr(\mathbf{M}, \mathbf{i}, \mathbf{e}) = \mathbf{m}$, i.e., for which \mathbf{m} is the most probable explanation to \mathbf{M} , and decide upon the majority. However, this algorithm trivially runs in exponential time, even if the treewidth of the network is bounded: there are exponentially many joint value assignments to $\mathbf{M} \cup \mathbf{I}$. We present a randomized algorithm that performs much better in many circumstances.

```

for  $n = 1$  to  $N$  do
  Choose  $\mathbf{i}$  at random
  Determine  $\mathbf{m} = \text{argmax}_\mathbf{M} \Pr(\mathbf{M}, \mathbf{i}, \mathbf{e})$ 
  Count the joint value assignments  $(\mathbf{m}, \mathbf{i})$ 
end for
Decide upon the majority and output  $\mathbf{m}_{\text{maj}}$ 

```

This randomized algorithm repeatedly picks a joint value assignment to \mathbf{I} at random, determines the most probable explanation, and at the end decides upon

the majority. Due to its stochastic nature, this algorithm is not guaranteed to give correct answers all the time. However, the error margin ϵ can be made sufficiently low by choosing N large enough, where the threshold value of N can be computed using the *Chernoff bound*: $N \geq \frac{1}{(p-\frac{1}{2})^2} \ln \frac{1}{\sqrt{\epsilon}}$. Assume we require an error margin of less than 0.1, then the number of repeats depends on the probability of picking a joint value assignment \mathbf{i} for which \mathbf{m}_{maj} is the most probable explanation⁷. If this probability is high (say $p = 0.85$), then N can be fairly low ($N \geq 10$), however, if the probability distribution is almost uniform then an exponential number of repetitions is needed.

If the majority is bounded (i.e., larger than a particular fixed threshold) we thus need only polynomially many repetitions to obtain any constant error rate, and thus the MOST SIMPLE EXPLANATION problem is in NP^{BPP} . When determining the most probable explanation is easy—in particular, when the treewidth of \mathcal{B} is low—the algorithm thus runs in polynomial time. Since the treewidth of \mathcal{B} is independent of the choice of \mathbf{m} and \mathbf{i} , MOST SIMPLE EXPLANATION can be decided in polynomial time, with a small possibility of error, when the treewidth is low and the majority is sufficiently large. Note that this result is congruent with observations by Vul et. al [27] where it was found that humans typically only use a very limited number of samples (sometimes even just a *single* sample) to make decisions under uncertainty, if the utility of that decision is high, relative to alternative decisions.

This sampling approach is conceptually different from Monte-Carlo-like approximations of probabilistic inference (as might be used in an approximation algorithm for PARTIAL MAP), where the marginalization process is approximated by taking samples of the intermediate variables according to their prior probability distribution. We argue, in line with the *Decision by Sampling* model as discussed in the beginning of this section, that it may be plausible that one *does not marginalize at all* to obtain most probable explanations, but often bases ones judgment on what appears to hold for the vast majority of intermediate variables.

4 Most Informative Explanations

In this section we introduce Most Informative Explanations that enhance the traditional interpretation of Bayesian abduction with a notion of *information-richness*. This alternative interpretation is based on the intuitive idea that the explanation set (and thus the specificity of the most probable explanation) is not always fixed, but can be flexibly adjusted to be more general or more specific, depending on the probability of the most probable explanation for a particular size of the explanation set. Since we marginalize over variables outside both the evidence and the explanation set, the probability of the most probable explanation is by definition larger when the explanation set becomes

⁷For ease of exposure, we assume here that there are only two competing joint value assignments, and consequently, this probability is by definition larger than $\frac{1}{2}$. If this assumption is violated, the standard Chernoff bound can not be used to compute N for a given error margin.

smaller, however the information carried in that explanation is less specific, as was illustrated by the *Dr. Brown*-example in the introduction. In this section, we formalize this problem, analyze its computational complexity and show under what circumstances the problem becomes tractable.

4.1 Computational Complexity

The decision version of MOST INFORMATIVE EXPLANATION will be formulated as follows.

MOST INFORMATIVE EXPLANATION

Instance: A probabilistic network $\mathcal{B} = (\mathbf{G}, \Gamma)$, where \mathbf{V} is partitioned into a set of evidence nodes \mathbf{E} with a joint value assignment \mathbf{e} , an explanation set \mathbf{M} , and intermediate variables \mathbf{I} ; a rational number $0 \leq q < 1$ and a natural number $0 \leq l \leq |\mathbf{M}|$.

Question: Is there a joint value assignment \mathbf{m}_l to a subset of length l of the nodes in \mathbf{M} such that $\Pr(\mathbf{m}_l, \mathbf{e}) > q$?

With respect to the *Dr. Brown* example, the evidence \mathbf{e} would contain observations as “Coughs-with-phlegm = TRUE” and “Has-pain-while-breathing = TRUE”, the explanation set \mathbf{M} would contain the variables Ill (yes/no), Pneumonia (severe/moderate/absent) and Pneumonia-type (hospital-acquired/community-acquired/SARS/none), and the set of intermediate variables would contain variables like Hosts-Streptococcus-Pneumoniae and Has-Impaired-Immune-System. The MOST INFORMATIVE EXPLANATION problem, in its decision variant, asks whether there is an explanation for at least l variables whose posterior probability exceeds a threshold q .

Theorem 4.1. MOST INFORMATIVE EXPLANATION is NP^{PP} -complete.

Proof. Membership can be shown by non-deterministically guessing \mathbf{m}_l and deciding, using the PP oracle, whether $\Pr(\mathbf{m}_l, \mathbf{e}) > q$. NP^{PP} -hardness follows since MOST INFORMATIVE EXPLANATION has PARTIAL MAP as a special case: take $l = |\mathbf{M}|$. \square

If $l = 0$ then MOST INFORMATIVE EXPLANATION degenerates to INFERENCE. If $l = |\mathbf{M}|$ then MOST INFORMATIVE EXPLANATION degenerates to PARTIAL MAP, and if in addition $\mathbf{I} = \emptyset$ then MOST INFORMATIVE EXPLANATION degenerates to MPE. Furthermore, MOST INFORMATIVE EXPLANATION inherits the inapproximability results of PARTIAL MAP [13].

While PARTIAL MAP is fixed parameter tractable for $(1 - p, \text{tw})$, i.e., PARTIAL MAP can be solved fast when the treewidth of the network is bounded and the most probable explanation has a high probability, this may not hold for MOST INFORMATIVE EXPLANATION since we need to choose \mathbf{m}_l which by itself is a source of complexity. However, Bodlaenders algorithm for PARTIAL MAP, which is fixed parameter tractable for $(1 - p, \text{tw})$ [13], can be adjusted by branching on each l -sized subset of the explanation set. Since there are

$\binom{l}{|\mathbf{M}|}$ many such subsets, the number of subsets to consider is low when either l is low (the subset is small) or $|\mathbf{M}| - l$ is low (the subset is large), but increases with l approaching $\frac{|\mathbf{M}|}{2}$. Therefore, MOST INFORMATIVE EXPLANATION is fixed parameter tractable for both $(1 - p, \text{tw}, l)$ and $(1 - p, \text{tw}, |\mathbf{M}| - l)$. Since MOST INFORMATIVE EXPLANATION is a generalization of both PARTIAL MAP (for $l = |\mathbf{M}|$) and INFERENCE (for $l = 0$) it follows that MOST INFORMATIVE EXPLANATION remains intractable for the sets of parameters $(1 - p, l)$, $(1 - p, |\mathbf{M}| - l)$ and $(\text{tw}, |\mathbf{M}| - l)$, but is fixed parameter tractable for the set of parameters (tw, l) .

We do not know whether MOST INFORMATIVE EXPLANATION is fixed parameter tractable for the set of parameters $(1 - p, \text{tw})$. However, we show that a related problem variant, which we will denote as MINIMAL EVIDENCE SUPPORT, is fixed parameter *intractable* for the set of parameters $(1 - p, \text{tw})$ ⁸ using a reduction from VERTEX COVER: given an undirected graph $\mathbf{G} = (\mathbf{V}, \mathbf{E})$, does it have a vertex cover of size at most k ? MINIMAL EVIDENCE SUPPORT, is defined as follows:

MINIMAL EVIDENCE SUPPORT

Instance: As in MOST INFORMATIVE EXPLANATION.

Question: Is there a joint value assignment \mathbf{m}_l to a subset of length l of the nodes in \mathbf{M} such that $\Pr(\mathbf{e} | \mathbf{m}_l) > q$?

Note that this problem is related to MOST INFORMATIVE EXPLANATION as $\Pr(\mathbf{e} | \mathbf{m}_l) = \frac{\Pr(\mathbf{m}_l, \mathbf{e})}{\Pr(\mathbf{m}_l)}$. As a running example, we use the graph \mathbf{G}_{ex} as depicted in Figure 2; it is easy to see that this graph has a vertex cover of size 2 (e.g., $\{V_2, V_3\}$), but has no vertex cover of size 1.




Figure 2: Example of a VERTEX COVER instance \mathbf{G}_{ex} with four nodes and four edges. A possible vertex cover may be $\{V_1, V_2\}$ or $\{V_2, V_3\}$

We reduce VERTEX COVER to MINIMAL EVIDENCE SUPPORT, ensuring that the treewidth of the resulting network is bounded and the posterior probability

⁸To be precise: is Para-NP-hard for parameter set $(1 - p, \text{tw})$.

distribution is such that there is a joint value assignment \mathbf{m}_k to a subset of length at most k of a designated set \mathbf{M} with probability $\Pr(\mathbf{e} | \mathbf{m}_k) = 1$ if and only if the original graph has a vertex cover of size at most k . From a given VERTEX COVER instance (\mathbf{G}, k) with n variables and m edges we construct a MINIMAL EVIDENCE SUPPORT instance $(\mathcal{B}, \mathbf{M}, \mathbf{E}, \mathbf{I}, k)$ as follows. For every vertex $V_i, 1 \leq i \leq n$ in $\mathbf{V}(\mathbf{G})$, we construct a binary root node V_i in \mathcal{B} with values *Selected* and *Not-Selected* and uniform probability, and a binary node C_i in \mathcal{B} with values *Covered* and *Not-Covered*. Furthermore, we introduce a root node E in \mathcal{B} , with m values (one for each edge in the original graph), uniformly distributed. Each node C_i has as parents V_i and E ; furthermore, there are arcs from C_i to C_{i+1} for $i = 1, \dots, n-1$. The probability distribution $\Pr(C_1 | V_1, E)$ and for $\Pr(C_{i+1} | C_i, V_{i+1}, E), i = 1, \dots, n-1$, is as follows.

$$\Pr(C_i = \textit{Covered} | C_{i-1}, V_i, E) = \begin{cases} 1 & \text{if } i = 1, E = e_j, V_i = \textit{Selected}, \\ & \text{and } V_i \text{ covers } e_j \\ 1 & \text{if } i > 1, E = e_j, C_{i-1} = \textit{Not-Covered}, \\ & V_i = \textit{Selected} \text{ and } V_i \text{ covers } e_j \\ 1 & \text{if } i > 1 \text{ and } C_{i-1} = \textit{Covered} \\ 0 & \text{otherwise} \end{cases}$$

Figure 3 shows the thus constructed network for our example graph \mathbf{G}_{ex} , and a tree-decomposition of the moralization of \mathbf{G}_{ex} . We set $\mathbf{M} = \{V_1, \dots, V_n\}$ and we have evidence $\mathbf{E} = \{C_n = \textit{Covered}\}$. The intuition behind the construction is as follows. The nodes V_i are used to select vertices that are to be in the vertex cover, the node E is used to assure that all edges are to be covered. For every node C_i we have that $\Pr(C_i = \textit{Covered} | C_{i-1}, V_i, E) = 1$ if the selected edge is covered by selecting V_i or if the selected edge was already covered by $V_j, j < i$. In particular we have that, for any joint value assignment \mathbf{m}_k that sets k variables V_i to *Selected*, $\Pr(C_n = \textit{Covered} | \mathbf{m}_k) = \sum_{e_j} \Pr(C_n = \textit{Covered} | E = e_j, \mathbf{m}_k) \cdot \Pr(E = e_j) = 1$ if and only if \mathbf{m}_k corresponds to a vertex cover to \mathbf{G} , i.e., a set of k vertices that covers all edges in \mathbf{G} . If, on the other hand, there does not exist a vertex cover of k vertices (i.e., at most $m-1$ edges are covered with k vertices), then $\Pr(C_n = \textit{Covered} | \mathbf{m}_k)$ is at most $1 - \frac{1}{m} \cdot \frac{1}{2^{n-k}}$ as we marginalize over the $n-k$ variables in \mathbf{M} that are not selected. For any graph with at least one edge, this probability is at least $\frac{1}{2}$, but strictly smaller than 1.

Observe that the treewidth of the thus constructed network is bounded (in fact, is equal to 3) as illustrated by Figure 3.

Theorem 4.2. *MINIMAL EVIDENCE SUPPORT is Para-NP-hard for networks with bounded treewidth and bounded posterior probability.*

Proof. We reduce MINIMAL EVIDENCE SUPPORT from VERTEX COVER. Let (\mathbf{G}, k) be an instance of VERTEX COVER with n variables and m edges. From (\mathbf{G}, k) , we construct the MOST INFORMATIVE EXPLANATION instance $(\mathcal{B}, \mathbf{M}, \mathbf{E}, \mathbf{I}, k)$ as shown above. Let $\mathbf{M} = \{V_1, \dots, V_n\}$ and $\mathbf{E} = \{C_n\}$ with evidence $C_n = \textit{Covered}$. If (\mathbf{G}, k) is a positive instance of VERTEX COVER,




Figure 3: Example of construction of a probabilistic network $\mathcal{B}_{\phi_{\text{ex}}}$ for the VERTEX COVER instance \mathbf{G}_{ex} , and a tree-decomposition of its moralization

then \mathbf{G} has a vertex cover of size k and thus there exists a set $\mathbf{m}_{\mathbf{k}}$ such that $\Pr(C_n = \text{Covered} | \mathbf{m}_{\mathbf{k}}) = 1$. Vice versa, if $\Pr(C_n = \text{Covered} | \mathbf{m}_{\mathbf{k}}) = 1$, then the variables in $\mathbf{m}_{\mathbf{k}}$ form a vertex cover of size k . Thus, if we can decide whether $\Pr(C_n = \text{Covered} | \mathbf{m}_{\mathbf{k}}) > 1 - \frac{1}{m} \cdot \frac{1}{2^{n-k}}$, then we can decide VERTEX COVER as well. Note that the treewidth of \mathcal{B} is bounded and that $(\mathcal{B}, \mathbf{M}, \mathbf{E}, \mathbf{I}, k)$ can be constructed from (\mathbf{G}, k) in polynomial time, hence, MINIMAL EVIDENCE SUPPORT is NP-hard for networks with bounded treewidth and bounded posterior probability. \square

Observe that $\Pr(C_n = \text{Covered}, \mathbf{m}_{\mathbf{k}}) = \Pr(C_n = \text{Covered} | \mathbf{m}_{\mathbf{k}}) \cdot \Pr(\mathbf{m}_{\mathbf{k}}) = \frac{\Pr(C_n = \text{Covered} | \mathbf{m}_{\mathbf{k}})}{2^k}$ as the variables V_i are uniformly distributed and, for any subset of k variables, the prior probability that they are all set to TRUE is $\frac{1}{2^k}$. Thus, the proof construction used above cannot be used to prove NP-hardness of MOST INFORMATIVE EXPLANATION as the posterior probability is not bounded.

The fixed parameter tractability results for MOST INFORMATIVE EXPLANATION are summarized in Table 4.

5 Conclusion

In this paper we proposed two alternative notions of abduction in probabilistic networks, based on observations and intuitions in cognitive science. MOST SIMPLE EXPLANATION focuses on the problem of finding a *parsimonious* explanation, MOST SIMPLE EXPLANATION focuses on the problem of finding an *informative* explanation.

We proved NP^{PP} -hardness of both problem variants, placing them in the same complexity class as PARTIAL MAP, however, with different fixed parameter tractability results. In particular, MOST SIMPLE EXPLANATION is fixed parameter tractable when the treewidth of the network is bounded and there is

Parameters	Complexity	Remarks
$(1 - p, \text{tw}, l)$	FPT	Variant of the algorithm discussed in [13]
$(1 - p, \text{tw}, M - l)$	FPT	Variant of the algorithm discussed in [13]
$(1 - p, l)$	PP-hard	For $l = 0$ the problem reduces to INFERENCE parameterized by $1 - p$ which remains PP-hard [14]
$(1 - p, M - l)$	PP-hard	For $ M = l$ the problem reduces to PARTIAL MAP parameterized by $1 - p$ which remains PP-hard [13]
(tw, l)	FPT	for $l = 0$ the problem reduces to INFERENCE which is FPT for treewidth [16]
$(\text{tw}, M - l)$	NP-hard	For $ M = l$ the problem reduces to PARTIAL MAP parameterized by tw which remains NP-hard [20]
$(\text{tw}, 1 - p)$	NP-hard	For the related problem variant MINIMAL EVIDENCE SUPPORT, using a reduction from VERTEX COVER (Thm. 4.2)
tw	NP-hard	Because of NP-hardness of the problem parameterized by $(\text{tw}, M - l)$
$1 - p$	PP-hard	Because of PP-hardness of the problem parameterized by $(1 - p, l)$
l	PP-hard	For $l = 0$ the problem reduces to INFERENCE which is PP-hard [17]
$ M - l$	NP^{PP} -hard	For $ M = l$ the problem reduces to PARTIAL MAP which is NP^{PP} -hard [20]

Table 4: Summary of FPT-results for MOST INFORMATIVE EXPLANATION

a bounded majority of the joint value assignments to \mathbf{I} for which a particular explanation \mathbf{m} is the most probable explanation. Informally, one can envisage some sort of sampling: if \mathbf{m} is the most probable explanation for almost all joint value assignments to the set of intermediate variables \mathbf{I} , then only a few samples will suffice to say so with a relatively healthy error margin⁹.

We see MOST INFORMATIVE EXPLANATION as an extension, rather than an

⁹Note that this is an explanation at Marr’s [18] *algorithmic*, rather than *computational*, level since there may be alternative fast algorithms that do not use sampling at all.

alternative, to PARTIAL MAP. The formulation of the problem variant with a flexible size of the explanation set may serve as a starting point for more thorough investigations regarding the computational model of the cognitive task of deciding *what* to explain.

We do not put forward claims on whether MOST SIMPLE EXPLANATION or PARTIAL MAP (or yet an other notion of abduction) better represent the cognitive processes that are modeled by abductive Bayesian reasoning. One highly speculative hypothesis might be that a ‘sampling’ approach (i.e., MOST SIMPLE EXPLANATION) is used when the set of intermediate variables \mathbf{I} does not contain variables whose values may have a non-trivial impact on the most probable explanation, and a ‘marginalization’ approach (i.e., PARTIAL MAP) is used when they do. Maybe a combination of the two problems is used: marginalize over the relevant variables and sample the irrelevant ones (but which are the relevant variables?). Further elaboration, however, is beyond the scope of this paper.

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, Pim Haselager, Todd Wareham, Moritz Müller and Hans Bodlaender for valuable discussion regarding various ideas put forward in this paper and comments on earlier drafts.

References

- [1] C. L. Baker, R. Saxe, and J. B. Tenenbaum. Action understanding as inverse planning. *Cognition*, 113(3):329–349, 2009.
- [2] H. L. Bodlaender. Treewidth: Algorithmic techniques and results. In *Proceedings of the 22nd International Symposium on Mathematical Foundations of Computer Science*, number 1295 in Lecture Notes in Computer Science, pages 19–36, Berlin, 1997. Springer-Verlag.
- [3] H. L. Bodlaender, F. van den Eijkhof, and L. C. van der Gaag. On the complexity of the MPA problem in probabilistic networks. In F. van Harmelen, editor, *Proceedings of the 15th European Conference on Artificial Intelligence*, pages 675–679, 2002.
- [4] R. H. Cuijpers, H. T. van Schie, M. Koppen, W. Erlhagen, and H. Bekkering. Neural networks. *Goals and means in action observation: A computational approach*, 19:311–322, 2006.
- [5] R. G. Downey and M. R. Fellows. *Parameterized Complexity*. Springer Verlag, Berlin, 1999.

- [6] K. T. Fann. *Peirce's theory of abduction*. Springer, Berlin, 1970.
- [7] M. R. Garey and D. S. Johnson. *Computers and Intractability. A Guide to the Theory of NP-Completeness*. W. H. Freeman and Co., San Francisco, CA, 1979.
- [8] J. T. Gill. Computational complexity of Probabilistic Turing Machines. *SIAM Journal of Computing*, 6(4):675–695, 1977.
- [9] D. H. Glass. Coherence measures and inference to the best explanation. *Synthese*, 157:275–296, 2007.
- [10] D. H. Glass. Inference to the best explanation: does it track truth? *Synthese*, page In press, 2010.
- [11] C. G. Hempel. *Aspects of Scientific Explanation*. Free Press, New York, 1965.
- [12] F. V. Jensen and T. D Nielsen. *Bayesian Networks and Decision Graphs*. Springer Verlag, New York, NY, second edition, 2007.
- [13] J. Kwisthout. Most Probable Explanations in Bayesian networks: complexity and tractability. Technical Report ICIS–R10001, Radboud University Nijmegen, 2010.
- [14] J. Kwisthout. On the complexity of probabilistic inference. Technical report, Radboud University Nijmegen, In Preparation.
- [15] C. Lacave and F. J. Díez. A review of explanation methods for Bayesian networks. *The Knowledge Engineering Review*, 17(2):107–127, 2002.
- [16] S. L. Lauritzen and D. J. Spiegelhalter. Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society*, 50(2):157–224, 1988.
- [17] M. L. Littman, J. Goldsmith, and M. Mundhenk. The computational complexity of probabilistic planning. *Journal of Artificial Intelligence Research*, 9:1–36, 1998.
- [18] D. Marr. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Freeman, New York, NY, 1982.
- [19] D. Nilsson. An efficient algorithm for finding the M most probable configurations in probabilistic expert systems. *Statistics and Computing*, 8:159–173, 1998.
- [20] J. D. Park and A. Darwiche. Complexity results and approximation settings for MAP explanations. *Journal of Artificial Intelligence Research*, 21:101–133, 2004.

- [21] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, Palo Alto, CA, 1988.
- [22] J. A. Reggia and Y. Peng. Modeling diagnostic reasoning: a summary of parsimonious covering theory. *Computer methods and programs in biomedicine*, 25(2):125–134, 1987.
- [23] S. E. Shimony. Finding MAPs for belief networks is NP-hard. *Artificial Intelligence*, 68(2):399–410, 1994.
- [24] N. Stewart, N. Chater, and G. D. A. Brown. Decision by sampling. *Cognitive Psychology*, 53:1–26, 2006.
- [25] B. K. Sy. Reasoning MPE to multiply connected belief networks using message-passing. In P. Rosenbloom and P. Szolovits, editors, *Proceedings of the Tenth National Conference on Artificial Intelligence*, pages 570–576. AAAI Press, Arlington, Va, 1992.
- [26] I. van Rooij and T. Wareham. Parameterized complexity in cognitive modeling: Foundations, applications and opportunities. *The Computer Journal*, 51(3):385–404, 2008.
- [27] E. Vul, N. D. Goodman, T. L. Griffiths, and J. B. Tenenbaum. One and done? optimal decisions from very few samples. In *31st Annual Meeting of the Cognitive Science Society*, 2009.
- [28] C. Yuan, T. Lu, and M. Druzdzel. Annealed MAP. In J. Halpern, editor, *Proceedings of the Twentieth Conference Annual Conference on Uncertainty in Artificial Intelligence*, pages 628–635. AUAI Press, Arlington, Va, 2004.
- [29] A. Yuille and D. Kersten. Vision as Bayesian inference: analysis by synthesis? *TRENDS in Cognitive Sciences*, 10(7):301–308, 2006.