Supplementary material for:

How Action Understanding can be Rational, Bayesian *and* Tractable

by Mark Blokpoel, Johan Kwisthout, Theo P. van der Weide and Iris van Rooij

In this supplementary material we provide proofs of the complexity-theoretic results we claim in the main paper. First we provide an overview of the concepts and definitions from Bayesian modeling (Section 1) and complexity theory (Section 2). Secondly we introduce a general form of the original Bayesian Inverse Planning model (Section 3). And finally we verify the statements we made in our text about the complexity of the models (Section 4).

1 Preliminaries from Bayesian modeling

For readers unfamiliar with basic notations from Bayesian modeling we review some of the basics relevant for our purpose. For details we refer the reader to the sources in the text.

A *Bayesian network (BN)* (Pearl, 1988; Ghahramani, 1998; Jensen & Nielsen, 2007) is a tuple denoted by $\mathcal{B} = (\mathbf{G}, \Gamma)$, where $\mathbf{G}$ is a directed acyclic graph $\mathbf{G} = (\mathbf{V}, \mathbf{A})$ that models the stochastic variables and their dependencies and $\Gamma = \{P_X | X \in \mathbf{V}\}$ is the set of conditional probability distributions $P(X \mid \mathbf{y})$ for each joint value assignment $\mathbf{y}$ to the parents of $X \in \mathbf{G}$. For clarity a BN is usualy depicted by a graph, where directed edges $(X, Y) \in \mathbf{A}$ represent dependencies $P(Y \mid X)$.

Let $\mathbf{W}$ be a set of variables. In a BN a *joint value assignment* $\mathbf{w}$ for $\mathbf{W}$ is an adjustment to the prior probabilities for each variable $V_i \in \mathbf{V}$ and each associated value $w_i \in \mathbf{w}$ such that $P(W_i = w_i) = 1$ and $P(W_i \neq w_i) = 0$. When a joint value assignment is observed or known, it is often called *evidence* $\mathbf{e}$ for a particular set of variables $\mathbf{E} \subseteq \mathbf{V}$.

A *joint probability distribution* for a set of variables $\mathbf{W}$ defines all the probabilities of all combinations of values for the variables in $\mathbf{W}$. Formally let ξ denote a Boolean algebra of propositions spanned by $\mathbf{V}$. The function $P : \xi \rightarrow [0, 1]$ is a joint probability distribution on $\mathbf{V}$ if the following conditions hold:

- $0 \leq P(a) \leq 1$, for all $a \in \xi$;
- $P(TRUE) = 1$;
- $P(FALSE) = 0$;
- for all $a, b \in \xi$, if $a \wedge b \equiv FALSE$ then $P(a \vee b) = P(a) + P(b)$.

Dynamic BNs (dBN) (Ghahramani, 1998) are BNs that represent sequences of variables (called a slice), often related to time. Each slice is a BN $\mathcal{B}_t = (\mathbf{G}, \Gamma)$ with an index $t \in \mathbb{N}$. Let $I \subseteq \mathbf{V}$ be the set of input variables and $O \subseteq \mathbf{V}$ be the set of output variables such that $\forall_{t,t'}[I_t = I_{t'} \wedge O_t = O_{t'}]$ and $\forall_{t,i \in I} \exists_{o \in O}[P(i_{t+1} \mid o_t) \in \Gamma]$.

A common problem in Bayesian modeling is finding the Most Probable Explanation (MPE) for certain variables, denoted as the *evidence set*, given certain evidence. In fact, inverse Bayesian planning is a special case of MPE.

Most Probable Explanation
Input: A probabilistic network $\mathcal{B} = (\mathbf{G}, \Gamma)$, where $\mathbf{V}$ is partitioned into a set of evidence nodes $\mathbf{E}$ with a joint value assignment $\mathbf{e}$ and an explanation set $\mathbf{M}$, such that $\mathbf{E} \cup \mathbf{M} = \mathbf{V}$.
Output: What is the most probable joint value assignment $\mathbf{m}$ to the nodes in $\mathbf{M}$ given evidence $\mathbf{e}$?

The *tree-decomposition* (Robertson & Seymour, 1986) of any graph $\mathbf{G} = (\mathbf{V}, \mathbf{E})$ is a pair $< T, \mathcal{X} >$, where $T = (I, F)$ is a tree and $\mathcal{X} = \{\mathbf{X_i} | i \in I\}$ is a family of subsets (called bags) of $\mathbf{V}$, one for each node of T, such that:

- $\bigcup_{i \in I} \mathbf{X_i} = \mathbf{V}$,
- for every edge $(V, W) \in \mathbf{E}$ there exists an $i \in I$ with $V \in \mathbf{X_i}$ and $W \in \mathbf{X_i}$,
- for every $i, j, k \in I$: if j is on the path from i to k in T, then $\mathbf{X_i} \cap \mathbf{X_k} \subseteq \mathbf{X_j}$.

Treewidth (Robertson & Seymour, 1986) of a BN $\mathcal{B}$ is defined as the minimum width over all tree-decompositions of the moralized graph of $\mathcal{B}$. The width of a tree-decomposition $((I, F), \{\mathbf{X_i} | i \in I\})$ is $\max\limits_{i \in I} |\mathbf{X_i}| - 1$.

2 Preliminaries from Complexity theory

We also assume the reader is familiar with basic notations from complexity theory (Garey & Johnson, 1979) but review some basics of parameterized complexity theory (Downey, Fellows, & Langston, 2008) relevant for our purpose. For details we refer the reader to textbooks by Garey and Johnson and Downey et al..

Let $P : I \rightarrow O$ be a problem with input parameters $k_1, k_2, \ldots, k_m$. Then P is *fixed-parameter tractable* for parameter set $K = \{k_1, k_2, \ldots, k_m\}$ if there exists at least one algorithm that computes P for any input of size n in time $O(f(k_1, k_2, \ldots, k_m) n^c)$, where f is an arbitrary computable function and c is a constant. If no such algorithm exists then P is said to be *fixed-parameter (fp-) intractable* for K.

Further more let P and Q be problems where Q is a special case of P. Then if P is fp-tractable, Q is also fp-tractable.

If a problem P is fp-intractable for a parameter set K, than P is fp-intractable for any subset $K' \subseteq K$.

3 Preliminaries from Inverse Bayesian Planning

Finally we assume the reader is familiar with the inverse Baysian planning (BIP) theory by Baker, Saxe, and Tenenbaum (2009). However we introduce a general framework of the theory and we define other special-cases based on this framework.

A BIP-Bayesian network (BIPBN) is a BN framework that we can use to define special cases such as M1, M2 and M3 by Baker et al. A BIPBN is a dBN $\mathcal{D}$ where each slice consists of a state variable $\mathbf{S_t} \in \mathbf{S}$ and action variable $\mathbf{A_t} \in \mathbf{A}$. Furthermore $\mathbf{G}$ can contain an arbitrary BN that encodes the goal(s). In this framework $P(\mathbf{A_t}) = P(\mathbf{A_t} \mid \mathbf{S_t}) P(\mathbf{S_t})$, $I = \mathbf{S_t}$, $O = \mathbf{A_t}$ and $P(\mathbf{S_{t+1}}) = P(\mathbf{S_{t+1}} \mid \mathbf{S_t}, \mathbf{A_t})$. All actions $\mathbf{A_t}$ are dependent on (at least one) goal in $\mathbf{G}$. Note that we did not include the world variable w and the noise variable β in our framework. These variables are constants and they are left out because any complexity results for this framework hold for the framework including w and β.

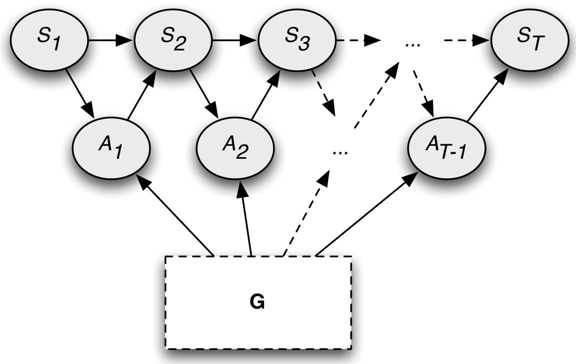

Figure 1: The BIPBN framework.

4 Complexity proofs

For all our proofs we assume that each variable $\mathbf{X}$ in a BN has a discrete and bounded set of values $\Omega(\mathbf{X})$, i.e. $|\Omega(\mathbf{X})| \leq k$, for a fixed $k \geq 0$.

4.1 M1, M2 and M3 are tractable

First we define M1, M2 and M3 as an input/output-problem. In this definition we assume the model's output is the most likely joint value assignment to $\mathbf{G}$ in the BIPBN (This makes M1, M2 and M3 special cases of MPE). Figure 2 contain graphical representations of M1, M2 and M3. We removed the variables κ and γ for reasons similair to why we leave out w and β (see Section 3).

Figure 2: BIP models M1, M2 and M3 and their tree-decompositions.

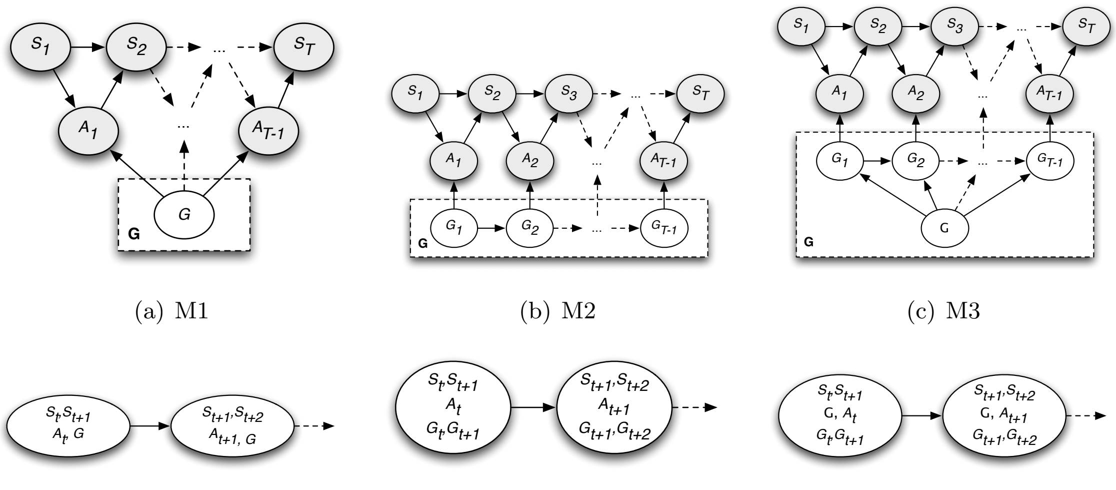

(a) M1 (b) M2 (c) M3

(d) M1 tree-decomposition (e) M2 tree-decomposition (f) M3 tree-decomposition

M1, M2 and M3
Input: A BIPBN $\mathcal{B} = (\mathbf{G}, \mathcal{D})$ and a joint value assignment (observations) $\mathbf{s}$ for $\mathbf{S}$ and $\mathbf{a}$ for $\mathbf{A}$. For M1, $\mathbf{G}$ contains one goal variable G and all actions are dependent on G; in M2 $\mathbf{G}$ contains a series of dependent goals $G_1, \ldots, G_{T-1}$ where G_t is dependent on G_{t-1} and each action A_t is dependent on G_t; in M3 $\mathbf{G}$ contains a series of dependent sub-goals $G_1, \ldots, G_{T-1}$ and a super-goal G where each sub-goal G_t is dependent on G and on G_{t-1} and each action A_t is dependent on G_t.
Output: The most likely joint value assignment to $\mathbf{G}$ given the evidence $\mathbf{s}$ and $\mathbf{a}$.

There are several algorithms (e.g. by Sy (1992) and Seroussi and Golmard (1994); see Kwisthout (2010) for an overview) that solve MPE in polynomial time when the treewidth of the moralized graph of $\mathcal{B}$ is bounded. More in particular, the runtime is $O(f(tw)g(p))$, where f is an exponential function based on the treewidth of $\mathcal{B}$ (tw) and g is a polynomial based on the number of cliques in $\mathcal{B}$ ($p \leq |\mathbf{V}|$). M1, M2 and M3 are special cases of MPE where the topology is restricted and treewidth is bounded.

The following results of treewidth are known, based on the tree-decompositions in Figure 2. These are not minimal but they are bounded by small numbers and thus suffice to prove M1, M2 and M3 tractable. Note that including the removed parameters β, γ, κ and w would increase the tree-width, but it would still be constant so the tractability result is also valid for the original model.

BIP model	*treewidth*
M1	4
M2	5
M3	6

Corollary 1. Because M1, M2 and M3 have treewidth ≤ 6, M1, M2 and M3 are tractable.

4.2 Proof multiple goals BIP is $\mathcal{NP}$-hard

Shimony (1994) proved finding MPE is $\mathcal{NP}$-hard in general BNs. We show that, even with its restricted topology, MULTIPLE GOALS BIP (MGBIP) is also $\mathcal{NP}$-hard. To prove MGBIP is $\mathcal{NP}$-hard, we provide a polynomial time reduction from DECISION-3SAT to DECISION-MGBIP and we argue that because DECISION-MGBIP is $\mathcal{NP}$-hard, MGBIP is also $\mathcal{NP}$-hard. First we need to define the decision variants of 3SAT and MGBIP.

DECISION 3SAT (D-3SAT)
Input: A tuple (U, C), where C is a set of clauses on Boolean variables U. Each clause is a disjunction of at most three variables.
Output: Does there exists a truth assignment to the variables in U that satisfies the conjunction of all clauses in C?

DECISION-MULTIPLE GOALS BIP (D-MGBIP)
Input: A BIPBN (see Figure 3) $\mathcal{B} = (\mathbf{G_1}, \ldots, \mathbf{G_k}, \mathcal{D})$ where , $k \geq 0$, and two sets of a joint value assignments (observations) $\mathbf{s}$ for $\mathbf{S}$ and $\mathbf{a}$ for $\mathbf{A}$. Furthermore, let $q \in [0, 1]$.
Output: Does there exist a joint value assignment $\mathbf{g}$ for $\mathbf{G}$ given evidence $\mathbf{s}$ and $\mathbf{a}$ such that $P(G = g) \geq q$?

Figure 3: The dynamic Bayesian Network that underlies MGBIP.

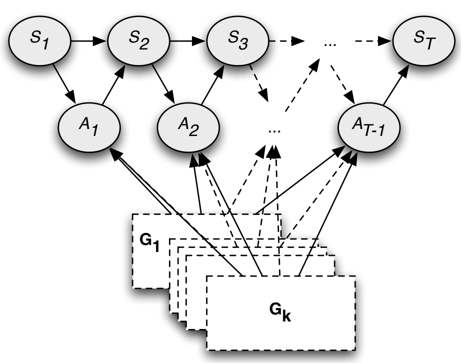

To rewrite a 3SAT instance to a D-MGBIP instance we represent a clause as an action variable in the BN. The conditional probability of the clause variable is constructed as:

Definition 1. *Clause variable probability distribution.* A clause variable is a node, that can model any clause of a 3SAT formula. A clause in 3SAT is the disjunction of at most three variables from the set $\{X_1, \ldots, X_k\}$, where each of the variables can be negated. The negations are encoded in the conditional probability of the clause. Let $\neg_p$ be true iff the p^{th} position of the clause is negated. We define the conditional probability of the clause variable as:

$$P(\mathbf{C} \mid X_h, X_i, X_j) = \begin{cases} 1 & (X_h \otimes \neg_1) \vee (X_i \otimes \neg_2) \vee (X_j \otimes \neg_3) \\ 0 & \text{otherwise} \end{cases}$$

Clause variable probability distribution for clauses with less variables can be defined analagous.

Lemma 1. D-MGBIP is $\mathcal{NP}$-hard.

The proof degrades dependencies in the BIPBN. Let C be dependend on A and B. Suppose we have to provide the conditional probabilities for the BN and each node can assume either *true* or *false*. Then we need to provide the following conditional probabilities:

$$P(C = true \mid A = true, B = true) = \alpha$$
$$P(C = true \mid A = true, B = false) = \beta$$
$$P(C = true \mid A = false, B = true) = \gamma$$
$$P(C = true \mid A = false, B = false) = \delta$$

If we set $\alpha = \beta$ and $\gamma = \delta$, then it does not matter what evidence we have for B. The conditional probability of $P(C \mid B)$ is the same, regardless of the value of B. In other words, C is not dependent on B. We will use this construction in the proof to degrade dependencies. Degraded dependencies will be denoted by dotted arrows in figures.

Proof. To reduce an instance of 3SAT φ to an instance of MGBIP $\mathcal{B}$, we create a multiple goal $\mathbf{G_i}$ containing one node G_i for each variable in φ. For each clause in φ an action with the corresponding clause probability distribution is created in $\mathcal{B}$ and for each conjunction in φ we create a conjunction node at state S_{t+1}, its conditional probability $P(S_{t+1} \mid S_t, A_t) = 1$ if $S_t = true$ and $A_t = true$ and 0 otherwise. Furthermore we define $S_0 = true$ in $\mathcal{B}$.

We degrade excess dependencies such that if there exists a valid truth assignment for the 3SAT-formula then there exists a joint value assignment $\mathbf{g}$ for $\mathbf{G_1}, \ldots, \mathbf{G_k}$ for which $P(\mathbf{g}) \geq q$. All dependencies between a goal node and a action node for which the *variable the goal node represents* is not present in the *clause the action node represents* are degraded. Furthermore, all dependencies between A_t and S_t are degraded. Figure 4 displays an example reduction with the degraded dependencies denoted as dotted arrows.

In $\mathcal{B}$ all state variables and actions variables are observed to be true and the prior probability distribution for each goal variable is normal.

The following conditions are met:

1. If φ is a yes-instance, then $\mathcal{B}$ is a yes-instance: For a 3SAT-formula to be satisfied, each clause must be satisfied. Per Definition 1 each action variable in $\mathcal{B}$ is *true* if and only if its corresponding clause is true. The probability of any joint value assignment $\mathbf{g}$ for $\mathbf{G_1}, \ldots, \mathbf{G_k}$ is 0 if it does not satisfy all clauses, or 1 if it does.

2. If $\mathcal{B}$ is a yes-instance, then φ is a yes-instance: Given the conditional probability $P\left(S_{t+1} \mid S_t,\ A_t\right)$, $\mathbf{G_1}, \ldots, \mathbf{G_k}$ need to be consistent with each clause variable in the BN. If $\mathcal{B}$ is a yes-instance then $P(\mathbf{g}) = 1$ and the joint value assignment $\mathbf{g}$ for $\mathbf{G_1}, \ldots, \mathbf{G_k}$ is consistent with each clause variable. Per definition of the clause variable's conditional probability distribution value assignment $\mathbf{g}$ satisfies each clause in φ.

3. The reduction runs in polynomial time: For each element in the 3SAT-formula only one node is created and a number of dependencies linear to the number of operators.

□

Lemma 2. If D-MGBIP is $\mathcal{NP}$-hard then MGBIP is also $\mathcal{NP}$-hard.

Proof. Assume there exists a polynomial time algorithm that solves MGBIP (viz. it returns the most probable explanation $\mathbf{g}$ for $\mathbf{G_1}, \ldots, \mathbf{G_k}$). Then together with the observations $\mathbf{s}$ for $\mathbf{S}$ and $\mathbf{a}$ for $\mathbf{A}$ we can compute $P\left(\mathbf{g} \mid \mathbf{s}, \mathbf{a}\right)$ in polynomial time, and check if it is $\geq q$. With a polynomial time algorithm for mgBIP we can solve d-mgBIP in polynomial time. We proved that d-mgBIP is $\mathcal{NP}$-hard, thus we have an inconsistency and we reject that mgBIP is solvable in polynomial time. □

Corollary 2. MGBIP is $\mathcal{NP}$-hard, because D-MGBIP is $\mathcal{NP}$-hard.

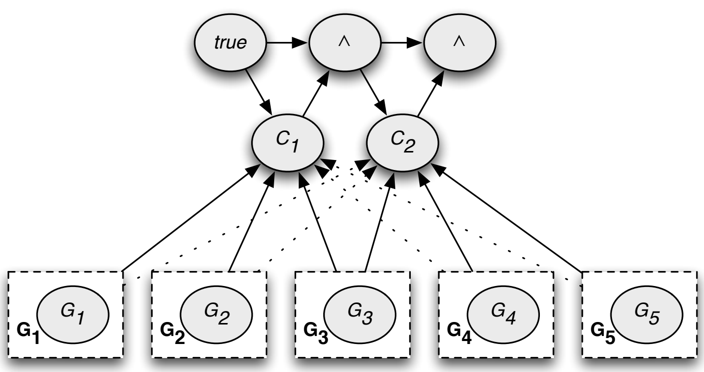

Figure 4: An example reduction from 3SAT to D-MGBIP. The clause $(G_1 \vee \neg G_2 \vee \neg G_3) \wedge (G_3 \vee G_4 \vee \neg G_5)$ is rewritten as a BN in mBIP. The conditional probabilities of clause$_0$ and clause$_1$ are: $P\left(\mathrm{C}_0 = \mathit{true} \mid G_1 = \mathit{true} \vee G_2 = \mathit{false} \vee G_3 = \mathit{false}\right) = 1$ and 0 otherwise, $P\left(\mathrm{C}_1 = \mathit{true} \mid G_1 = \mathit{true} \vee G_2 = \mathit{true} \vee G_3 = \mathit{false}\right) = 1$ and 0 otherwise.

4.3 Fixed-parameter tractability results for multiple goals BIP

Parameters s, a and g are, respectively, the maximum number of values per state, action and goal variable.

Corollary 3. MGBIP is not fixed-parameter tractable for $\{s, a, g\}$.

Proof. The $\mathcal{NP}$-hardness proof of MGBIP only uses a maximum of two values per variable (*true* or *false*), thus MGBIP is fp-intractable even when the number of values per variable is small. □

Parameter T is the maximum available observations.

Corollary 4. MGBIP is fp-intractable for $\{T\}$.

Proof. Even when the length of the observation is 1, with any number of multiple goals we can encode the entire 3SAT-formula in one action variable and a reduction from D-3SAT to D-MGBIP would be possible. Thus MGBIP is fp-intractable even when the maximum number of available observations is either small. □

Parameter $1/T$ is the maximum poverty of observations.

Corollary 5. MGBIP is fp-intractable for $\{1/T\}$.

Proof. If the reduction from 3SAT to D-MGBIP does not produce an instance with a large number of states such that $1/T$ is small, then we can add dummy state S'_t and action A'_t nodes. The conditional probability $P(S'_t \mid S_{t-1} = true, A_{t-1} = true) = 1$ or 0 otherwise and the conditional probability $P(A'_t \mid G_i = g_i, \ldots, G_j = g_j, S_{t-1} = true) = 1$ or 0 otherwise, where $g_i \ldots g_j$ can be any value (i.e. A'_t is independent of all goals). This means we can reduce any 3SAT instance to D-MGBIP while $1/T$ is small. □

Corollary 6. MGBIP is fp-intractable for $\{T, 1/T\}$.

Proof. Assume there exists an algorithm A that solves MGBIP in polynomial time, given T and $1/T$ are constant. This means we can solve MGBIP in polynomial time, given either T or $1/T$ is constant. This contradicts Corrolary 4 and Corrolary 5, thus we can conclude that such an algorithm does not exist. □

Because the above proofs do not assume more than two values for each variable we observe:

Result 1. MGBIP is fp-intractable for every subset of parameters $K \subseteq \{T, 1/T, g\}$.

Parameter k is the maximum number of multiple goals.

Proposition 1. MGBIP is fp-tractable for $\{k\}$.

Proof. We know that MPE is fixed-parameter tractable for treewidth (Kwisthout, 2010) and MG-BIP is a special case of MPE. Thus MGBIP is fixed-parameter tractable for treewidth. The treewidth of the BN underlying MGBIP grows as the number of goals increase (i.e. as the size of the input increases). Because treewidth is the only source of intractability for MGBIP and the number of goals is the only source that increases the treewidth we postulate MGBIP is fixed-parameter tractable for the number of multiple goals. □

Result 2. MGBIP is fp-tractable for parameter $\{k\}$.

Parameter $1 - p$ is the distance from complete certainty. Here p is the probability of the most probable explanation in MGBIP.

Proposition 2. MGBIP is fp-tractable $\{1 - p\}$.

Proof. It is known that MPE is fixed-parameter tractable for probability (Bodlaender, van den Eijkhof, & van der Gaag, 2002), in the sense that MPE can be solved efficiently if the probability of the most probable explanation is high. Given that MGBIP is a special case of MPE, MGBIP is fixed-parameter tractable for probability. □

Result 3. MGBIP is fp-tractable for parameter $\{1 - p\}$.

References

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

Bodlaender, H. L., van den Eijkhof, F., & van der Gaag, L. C. (2002). On the complexity of the MPA problem in probabilistic networks. In F. van Harmelen (Ed.), *Proceedings 15th european conference on artificial intelligence* (pp. 675–679). Amsterdam, the Netherlands: IOS Press.

Downey, R., Fellows, M., & Langston, M. (2008, Jan). The computer journal special issue on parameterized complexity. *The Computer Journal*.

Garey, M., & Johnson, D. (1979). Computers and intractability: A guide to the theory of NP-completeness. *orton.catie.ac.cr*.

Ghahramani, Z. (1998). Learning dynamic Bayesian networks. In *Adaptive processing of sequences and data structures* (pp. 168–197). Springer-Verlag.

Jensen, F. V., & Nielsen, T. D. (2007). *Bayesian networks and decision graphs* (Second ed.). New York: Springer Verlag.

Kwisthout, J. (2010, March). *Most Probable Explanations in Bayesian Networks: complexity and tractability* (Tech. Rep. Nos. ICIS–R10001). Radboud University Nijmegen.

Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference.* Morgan Kaufmann, Palo Alto.

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

Seroussi, B., & Golmard, J. (1994, Jan). An algorithm directly finding the k-most probable configurations in Bayesian networks. *International Journal of Approximate Reasoning*, *11*(3), 205–233.

Sy, B. K. (1992). Reasoning MPE to multiply connected belief networks using message-passing. In *Proceedings of the tenth national conference on artificial intelligence* (pp. 570–576).