

Bridging the gap between theory and practice of approximate Bayesian inference

Johan Kwisthout^{a*}, Iris van Rooij^a

^a Radboud University Nijmegen, Donders Institute for Brain, Cognition and Behaviour, Montessorilaan 3, 6525 HR Nijmegen, The Netherlands

*Corresponding author: tel. +31 24 3616288 email: j.kwisthout@donders.ru.nl

Abstract

In computational cognitive science, many cognitive processes seem to be successfully modeled as Bayesian computations. Yet, many such Bayesian computations have been proven to be computationally intractable (NP-hard) for unconstrained input domains, even if only an approximate solution is sought. This computational complexity result seems to be in strong contrast with the ease and speed with which humans can typically make the inferences that are modeled by Bayesian models. This contrast—between theory and practice—poses a considerable theoretical challenge for computational cognitive modelers: How can intractable Bayesian computations be transformed into computationally plausible ‘approximate’ models of human cognition? In this paper, three candidate notions of ‘approximation’ are discussed, each of which has been suggested in the cognitive science literature. We will sketch how (parameterized) computational complexity analyses can yield model variants that are tractable and which can serve as the basis of computationally plausible models of cognition.

Keywords: Bayesian inference; approximation; NP-hard; parameterized complexity theory; algorithms; computational explanation; cognitive modeling.

1. Introduction

Over the last decade, Bayesian modeling has become more and more important as a modeling framework in cognitive science. Many of such Bayesian models postulate that cognitive processes perform some form of Bayesian inference.¹ Examples of such models can be found in several different cognitive domains, including vision (Yuille & Kersten, 2006), language (Chater & Manning, 2006), decision making (Sloman & Hagmayer, 2006), motor planning (Wolpert & Ghahramani, 2005), eye movement control (Engbert & Krügel, 2010), and theory of mind (Baker, Saxe, & Tenenbaum, 2009; Cuijpers et al., 2006). These models often perform well at describing and predicting human behavior on small, well-structured experimental tasks. For instance, Bayesian models have been able to successfully model human inferences about an agent’s goal in a maze-like structure given the trajectory up so far (Baker et al., 2009) and predicting the most likely next action of a co-worker in a joint action task based on observed movements (Cuijpers et al., 2006). When confronted with task situations of real-world complexity, however, it is not evident if and how these models could scale up. The reason is that Bayesian computations are known to be computationally intractable (i.e., NP-hard) for unconstrained domains (Cooper, 1990; Shimony, 1994; Park & Darwiche, 2004; Kwisthout, 2009; Kwisthout, 2011). Informally, this means that the postulated computations would simply require an unrealistic amount of time for all but small inputs if they are to operate in unconstrained domains (van Rooij, 2008).

The known NP-hardness of (unconstrained) Bayesian models of cognition implies that such models cannot serve as computationally and psychologically plausible explanations of the tasks and processes which they model. Cognitive modelers using a Bayesian approach are thus confronted with a tractability paradox: humans seem capable of performing well in situations of real-world complexity in little time and with apparently little effort; yet, our best computational cognitive models consume an unrealistic amount of time or other cognitive resources for all but overly simplified toy domains. This paradox has been recognized by both opponents and proponents of Bayesian models, and has led to considerable debate about whether or not the human mind/brain performs such (rational) probabilistic computations (Chater, Oaksford, Nakisa, & Redington, 2003; Chater, Tenenbaum, & Yuille, 2006; Gigerenzer & Brighton, 2009; Gigerenzer, Hoffrage, & Goldstein, 2008; Kruschke, 2010; Oaksford & Chater, 2009; Sanborn et al., 2010).

A common response from proponents of Bayesian models of cognition has been that the tractability paradox is caused by a mistaken assumption, viz., the assumption that Bayesian models necessarily aim to be *exact* models of how human minds/brains form explanations. Instead, so these researchers argue, the models postulate computations that humans compute *approximately*. For instance, the mind may approximate Bayesian computations by employing heuristics, by sampling, or by striving for merely satisfactory solutions rather than optimal solutions (Chater et al., 2006; Sanborn et al., 2010).

Notwithstanding the appeal and apparent plausibility of this perspective, this standard response can at best be only part of the solution to the intractability paradox (Kwisthout, Wareham, & van Rooij, 2011). The reason is that most—if not all—forms of approximating Bayesian inference are *for unconstrained input domains* as computationally intractable as computing

¹ We will use the term ‘Bayesian inference’ to denote all sorts of computations using Bayesian models, like computing posterior probabilities and finding most probable explanations, as is common in the literature; we do not restrict this term to refer to the formal INFERENCE problem in Bayesian networks.

Bayesian inference exactly. For instance, even though empirical results suggest that in certain specific situations a few samples are sufficient to draw reasonable conclusions from a probability distribution (Vul, Goodman, Griffiths, & Tenenbaum, 2009); making decisions based on a ‘majority vote’ using samples remains provably intractable for general input domains (Kwisthout, 2010). Similarly, while approximate inference algorithms sometimes can perform well in practice (Park, 2002), approximating or “satisficing” Bayesian inference remains intractable for general domains (Abdelbar & Hedetniemi, 1998; Kwisthout, 2011; Park et al., 2004; Dagum & Luby, 1993).

The upshot of such negative computational complexity results for approximating Bayesian inference is that the assumption of ‘approximation’ by itself is insufficient to explain how Bayesian models can scale to situations of real-world complexity; and in effect, such models cannot yet claim computational plausibility as explanations of how humans make (Bayesian) inferences in everyday life (Kwisthout et al., 2011). Furthermore, by overlooking and not explicating the intrinsic complexity of approximating Bayesian inferences for certain situations, one may actually miss an opportunity to predict and explain under which circumstances humans inference can and does approximate optimal Bayesian inference and under which conditions it cannot and does not.

An alternative approach to tackle this intractability paradox is to study how the complexity of computations depends specifically on situational constraints. This approach draws among other things on the concepts and techniques from the mathematical theory of parameterized complexity theory (Downey & Fellows, 1999). Here, situational constraints are modeled by *parameters* $k_1, k_2 \dots k_m$, where each k is a property of the input. For instance, a Bayesian network may have many parameters, such as the number of cycles in the network, the maximum number of parents for any given node, the maximum length of a path between any two given nodes, the treewidth² of the network, etc.

It is known that certain NP-hard computations can be tractable for bounded ranges of such parameters. This is the case, for instance, if inferences can be computed in a time that grows non-polynomially only in the parameters and polynomial in the rest of the input size. In such situations the computations can be performed fast, even for very large inputs, provided only that the parameters take on small values in those large inputs. If similar parameterized tractability results can be obtained for Bayesian models of cognition, this means that the models predict that humans can make fast inferences in situations that are modeled by inputs with the same bounded parameter ranges.

This approach has previously been used to successfully identify situational constraints (modeled by parameter ranges) that yield tractability for both Bayesian and non-Bayesian models of cognition in a variety of domains, such as: analogy (van Rooij, Evans, Müller, Gedge, & Wareham, 2008), problem solving (Wareham, Evans, & van Rooij, 2011), similarity (Müller, van Rooij, & Wareham, 2009), action understanding (Blokpoel, Kwisthout, van der Weide, & van Rooij, 2010), and communication (van Rooij, Kwisthout, Blokpoel, Szymanik, Wareham, & Toni, 2011). For example, based on known parameterized complexity results on abduction in Bayesian networks (Bodlaender, van den Eijkhof, & Van der Gaag, 2002), van Rooij et al. (2011) derived that the inference made by an addressee to understand a communicator’s intention—as modeled by a Bayesian model—is tractable when the probability of the most probable goal is relatively high and the sender has few (instrumental) goals other than her communicative goals.

Until now, the parameterized complexity approach to dealing with intractability of cognitive models has focused on exact computations. Little research has been done, for instance, on the parameterized complexity of approximate Bayesian computations. Even the term ‘approximation’ appears to be often ill-defined. Take, for instance, again Bayesian abduction (computing the most probable explanation h out of a set of candidate explanations H , given a number of observations e , i.e., computing h such that $\Pr(H = h | e)$ is maximum). In the context of such a model ‘approximation’ could refer to finding explanations that are ‘almost as likely as’, ‘are similar to’, or ‘are likely to be’ the most probable explanation. Such ambiguity of the meaning of ‘approximation’ can lead to confusion and misunderstanding, as well as unsubstantiated claims of tractable approximability (see also van Rooij & Wareham (2012) for an extensive treatment of this issue for computational models of cognition in general). In this paper, we aim to explicate various notions of approximate Bayesian inference in an attempt to support a more thorough study of the approximability of Bayesian computations as featured in models of cognition. Additionally, we illustrate how an approximation approach and the parameterized complexity approach can be combined to identify the situational constraints that render Bayesian models tractably approximable.

The remainder of the paper is organized as follows. We will start by defining three distinct notions of approximation. Next, we will present a case study to illustrate the utility of the proposed combined approach for one of these three notions of approximation. The conclusion of our illustration will be that approximation is neither *panacea* nor *placebo*: it is not a ‘magical ingredient’ that makes intractable computations tractable *by default*. Yet, it may be a ‘necessary ingredient’ for achieving tractability under certain constraints. Finally, we conclude by discussing the broader implications of our case study and observations for dealing with the tractability paradox in Bayesian modeling in general.

² Treewidth is a graph-theoretical concept that can loosely be described as a measure on the “localness” of connections in a graph. We refer the reader to, for instance, Bodlaender (2006) for a formal definition and more details.

2. Notions of approximation

In everyday language, the concept ‘roughly’ is widely used: The ratio between the circumference and diameter of a circle is *roughly* 3.14; the Sun is *roughly* 100 times as large (and also 100 times as distant from the Earth) as the Moon; there is a mountain range on Mars that looks *roughly* like a human face; *roughly* speaking, height determines shoe size. We save a lot of effort if we don’t compute 98×102 exactly, but compute it as ‘roughly $100 \times 100 = 10,000$.’ When we come to see *approximation* as a synonym to *roughly computed*, in the above-sketched informal sense, the term has a pleasant association with ‘fast’, ‘easy’, and ‘without much effort’.

In computer science, however, there is a great variety on how fast and how well approximation algorithms perform for problems that are intractable to compute exactly. For example, the *knapsack* problem—in which one is to select a subset of numbers from a set of numbers so that the subset adds up to some pre-specified value—can be approximated within any desired guaranteed error margin in an amount of time that is polynomially bounded both in the input size and the error (Kellerer & Pferschy, 1999). In contrast, the *maximum clique* problem cannot be approximated at all in polynomial time, under the common assumption that $P \neq NP$ (Zuckerman, 2006). In these classic results in computer science, approximations are defined as solutions that have a *value* that is close to the optimal solution, i.e., the approximated value is within a given ratio of the value of the optimal solution. However, an approximation may also be defined as a solution that has a *structure* that is similar to the structure of the optimal solution, regardless of its value (Kwisthout, 2012; van Rooij & Wareham, 2012). The important difference between these forms of approximation can be understood as follows: a value-approximation of the *Bayesian abduction problem* is an explanation (i.e., a value assignment to H) that is only slightly less probable than the most probable solution, but need not resemble that optimal solution. Conversely, a structure-approximation is an explanation that *looks a lot* like the most probable explanation, but its probability may be much lower and need not even be close to optimal. Lastly, we may see approximations as solutions that are *likely to be* the optimal solution, but allow for a small margin of *expected* error. This type of approximation is common for probabilistic algorithms that are very useful for solving a certain class of problems (e.g., primality testing), as they are much faster than deterministic algorithms yet allow the possibility of error (Motwani & Raghavan, 1995). We will denote the latter notion of approximation as *expectation-approximation*. Note that a ‘good’ value-approximation does not imply a ‘good’ structure- or expectation-approximation, or vice versa (see Figure 1).

FIGURE 1 ABOUT HERE

When invoking approximation as an explanatory tool, cognitive modelers may utilize many (combinations) of these notions of approximation. For example, models may postulate that cognitive processes compute an approximation that resembles the optimal solution for a Bayesian inference problem as well as have an almost-as-high probability value (Chater & Oaksford, 1999). Combinations of value-approximation and expectation-approximation are also common in, e.g., the well-known PAC-learning framework for machine learning (Valiant, 1984). Consistent with the positions argued by Kwisthout et al. (2011, p. 780) and van Rooij and Wareham (2012), we submit that any claim that Bayesian models are tractable due to approximation should be supported by: (a) a precise definition of the notion of approximation the modeler assumes to be used; (b) in the case that the (formal) approximation problem in itself is NP-hard, a set of *problem parameters* that the modeler believes to be constrained in real-world situations (where humans perform the task well); (c) a formal proof that the chosen formal definition of the assumed approximation becomes *tractable* when the values of these parameters are so constrained; and (d) arguments or evidence that support the assumption in (b).

We will next illustrate with a case study how (a), (b), (c) and (d) can be implemented for Bayesian models using formal theoretical methods.

3. Case study: Most Simple Explanations

A common computational problem in Bayesian networks is inferring the best explanation out of a set of possible hypotheses, given some observations that are entered as evidence into the network. In such a network, the set of variables is typically partitioned into *explanation* variables (for which a value assignment is sought), *evidence* variables (whose value is observed), and *intermediate* variables (which are variables whose values are neither observed, nor part of the explanation set). This partition corresponds to the assumption that there are (many) variables that influence the resulting most probable hypothesis, but are not themselves observed. One way of dealing with such intermediate variables was proposed by Kwisthout (2010). Therein, the MOST SIMPLE EXPLANATION problem was defined as follows:

MOST SIMPLE EXPLANATION (informal)

Input: A probabilistic network, partitioned into a set of observed evidence nodes, an explanation set, and a set of intermediate variables.

Output: The joint value assignment to the explanation set that is the most probable explanation for the maximal number of instantiations of the intermediate variables.

Or more formally stated:

MOST SIMPLE EXPLANATION (formal)

Input: A probabilistic network B , partitioned into a set of evidence nodes E with a joint value assignment e , an explanation set H , and intermediate variables I .

Output: The joint value assignment h for which $\operatorname{argmax}_H \Pr(H, I = i, e) = h$ holds for the largest number of joint value i assignments to I .

Intuitively, one can think of MSE as the computational problem of finding the explanation that is most probable in the *majority of possible worlds*. Solving the MSE problem is, like many other computational problems defined over Bayesian networks, known to be intractable.⁵ Given this unfavorable complexity result, can MSE still serve as an adequate approximate model of abduction by humans in real-world situations? Using the four requirements introduced in the previous section we will show that the answer is *yes*: one can define an expectation-approximation algorithm for MSE that is tractable under conditions that seem to be met for human abduction.

a) *Give a precise definition of the notion of approximation that is used.* We use the following expectation-approximation to MSE, i.e., rather than exactly computing the solution to MSE, we want a solution that is very likely to be the MSE, but allows for a (guaranteed to be) small probability of error. In principle, such an expectation-approximation could be computed for MSE using a sampling algorithm that samples N random joint value assignments to the set of intermediate variables I and casts a majority vote over the samples. Moreover, using the so-called *Chernoff bound* (Chernoff, 1952) it can be exactly computed how many samples are required to expectation-approximate MSE to a given threshold probability. However, it is not possible to *tractably* expectation-approximate MSE, because the number of samples that are needed for a given degree of expectation-approximation can be exponential in the size of I .

b) *As the chosen approximation is intractable in general, we define problem parameters that we hypothesize to be constrained in the ‘real world’ situations that the model should capture.* A closer look at the MSE problem will reveal many problem parameters on the probability distribution and on the structure of the network. We consider three such parameters more closely:

1. The *treewidth* of a Bayesian network is a measure on the network structure. A formal definition of treewidth is mathematically non-trivial (see, e.g., Bodlaender, 2006, for details), but for our purposes it suffices to say that typically, if treewidth is low, the connections tend to be fairly *local* in the network. See Figure 2 for an illustrative example of this property.
2. The *cardinality* of a Bayesian network indicates the maximum number of values any variable can take; in binary variables the cardinality is two, but in principle variables can take many values.
3. The *relevancy* of the set of intermediate variables is a measure on the probability distribution associated with the network, and is defined as the probability that two random samples i_1 and i_2 would yield *different* most probable explanations. Informally, the intermediate variables have a low relevancy when there are only few possible worlds in which the most probable explanation deviates from the most probable explanation in the majority of worlds.

FIGURE 2 ABOUT HERE

c) *Give a formal proof that expectation-approximation of MSE becomes tractable when the values of these parameters are constrained.* It can be shown that if both the cardinality and the treewidth of the network are small, and in addition the intermediate variables have a low relevancy, then having only a few samples already suffices to solve the MSE with a low margin of error. Consider the following algorithm from Kwisthout (2010):

```
for  $n = 1$  to  $N$  do  
    Choose  $i$  at random  
    Determine  $h = \operatorname{argmax}_H \Pr(H, i, e)$ 
```

⁵ In fact, it is NP^{PP}-hard and thus resides in the same complexity class as, e.g., the MAP (Park & Darwiche, 2004) and Parameter Tuning (Kwisthout & van der Gaag, 2008) problems for Bayesian Networks.

Count the joint value assignments (h, i)

end for

Decide upon the majority and output h_{maj}

Computing $\operatorname{argmax}_H \Pr(H, i, e)$ can be done in time $O(c^{tw} \cdot n)$, where c denotes the cardinality of the network, tw the treewidth of the network, and n the number of variables of the network (Shimony, 1994). Also, according to the Chernoff bound, the number of samples N depends—for a fixed error rate—only on the probability that a randomly chosen sample actually is one of these possible worlds for which h_{maj} is the most probable explanation (i.e., depends only on the relevancy measure R) such that fewer samples are needed if this probability is high. We conclude that, for small c , small tw , and R close to 0, this algorithm tractably expectation-approximates the MSE problem.

d) Support the hypothesized constraints in b) with arguments or empirical findings. The dominant factor in the running time of computing $\operatorname{argmax}_H \Pr(H, i, e)$ is the treewidth of the network, as it is in the exponent. There is evidence in the machine learning literature that bounding treewidth while learning Bayesian networks from data prevents overfitting of the network (see e.g. Koller & Friedman, 2009). As there is reason to believe human learning tends to prevent overfitting as well (Gigerenzer & Brighton, 2009; Chater & Vitányi, 2003), the same principles may yield bounded treewidth for human belief networks. A second important factor in the running time is the relevancy of intermediate variables, i.e., the number of samples required to compute the MSE with low probability of error. While we are unaware of studies that *empirically* confirm that relevancy for human belief networks is low (i.e., number of required samples is small), there is independent theoretical research (e.g., Druzdzal, 1994; Stewart, Chater, & Brown, 2006; Vul et al., 2009) pointing in this direction.

4. Conclusion

Bayesian models of cognition are evidently successful in describing and predicting many cognitive phenomena. Despite this success, one important theoretical challenge remains to date: The tractability paradox, i.e., the apparent conflict between the computational intractability of Bayesian models for general input domains on the one hand, and the ease and speed with which humans perform the modeled tasks in situations of real-world complexity on the other. A common attempt at solving this paradox is to claim that minds/brains approximate, rather than exactly compute, the modeled inference problems. While we agree that approximation may be part of a solution, claims of ‘approximability’ alone are in the end explanatorily unsatisfactory as they are typically too ill-defined and sometimes even provably contradicted by known complexity results (Kwisthout et al., 2011).

In this paper we presented a method for solving the tractability paradox. Importantly, our method does not solve in one go the paradox for all Bayesian models. This is currently not possible, and may even be impossible in principle. Instead, our method describes how conditions for tractability can be studied and identified for different Bayesian models case by case. Crucial steps in our method are the following. First, one needs to decide on a relevant formal notion of approximation for the respective model. Here, we have discussed three such possible formal notions—viz., value-, structure-, and expectation-approximation—but altogether different notions or any combination of these are possible to adopt in our methodology as well. Second, using techniques from computational complexity theory one investigates under which conditions the respective model can be tractably approximated in the respective sense. Last, if there are good reasons to believe that the identified tractability conditions are met in those situations where humans can perform the modeled inferences quickly, the analysis has solved the tractability paradox for this specific Bayesian model.

We illustrated the use of our methodology for a specific Bayesian model: MSE (which stands for ‘Most Simple Explanation’). We found that MSE is intractable to expectation-approximate for unconstrained Bayesian networks, but is tractable to expectation-approximate for networks of bounded treewidth and with few relevant intermediate variables. Importantly, MSE is intractable to compute exactly under those same conditions. This set of results underscores our idea that approximation is not a panacea (approximation is not enough to yield tractability for all domains), but it is not a placebo either (it plays a necessary role in yielding tractability in certain conditions). We believe that by systematically studying this interplay between approximation and constraints on input domains for more Bayesian models, these models can achieve higher levels of computational plausibility, and at the same time serve as computational-level explanations of when and why humans are (and are not) good at making the rational inferences postulated by these models.

References

- Abdelbar, A.M., & Hedetniemi, S.M. (1998). Approximating MAPs for belief networks is NP-hard and other theorems. *Artificial Intelligence*, 102, 21–38.
- Baker, C.L., Saxe, R., & Tenenbaum, J.B. (2009). Action understanding as inverse planning. *Cognition*, 113, 329–349.
- Blokpoel, M., Kwisthout, J., van der Weide, T., & van Rooij, I. (2010). How action understanding can be rational, Bayesian and tractable. *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*, pp. 1643–1648.
- Bodlaender, H.L. (2006). Treewidth: characterizations, applications, and computations. *Proceedings of the 32nd International Workshop on Graph-Theoretic Concepts in Computer Science*, pp. 1–14.

- Bodlaender, H.L., van den Eijkhof, F., & van der Gaag, L.C. (2002). On the complexity of the MPA problem in probabilistic networks. *Proceedings of the 15th European Conference on Artificial Intelligence*, pp. 675–679.
- Chater, N., & Manning, C.D. (2006). Probabilistic models of language processing and acquisition. *TRENDS in Cognitive Sciences*, 10(7), 335–344.
- Chater, N., & Oaksford, M. (1999). Ten years of the rational analysis of cognition. *Trends in Cognitive Sciences*, 3, 57–65.
- Chater, N., Oaksford, M., Nakisa, R., & Redington, M. (2003). Fast, frugal and rational: How rational norms explain behavior. *Organizational Behavior And Human Decision Processes*, 90, 63–86.
- Chater, N., Tenenbaum, J., & Yuille, A. (2006). Probabilistic models of cognition: Conceptual foundations. *Trends In Cognitive Science*, 10((7)), 287–291.
- Chater, N., & Vitanyi, P. (2003). Simplicity: A unifying principle in cognitive science? *Trends in Cognitive Sciences*, 7(1), 19–22.
- Chernoff, H. (1952). A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Annals of Mathematical Statistics*, 23(4), 493–507.
- Cooper, G.F. (1990). The computational complexity of probabilistic inference using Bayesian belief networks. *Artificial Intelligence*, 42(2), 393–405.
- Cuijpers, R., Van Schie, H.T., Koppen, M., Erlhagen, W., & Bekkering, H. (2006). Goals and means in action observation: A computational approach. *Neural Networks*, 19, 311–322.
- Dagum, P., & Luby, M. (1993). Approximating probabilistic inference in Bayesian belief networks is NP-hard. *Artificial Intelligence*, 60 (1993), 141–153.
- Downey, R., & Fellows, M. (1999). *Parameterized Complexity*. Berlin: Springer.
- Druzdzel, M.J. (1994). Some properties of joint probability distributions. *Proceedings of the 10th Annual Conference on Uncertainty in Artificial Intelligence*, pp. 187–194.
- Engbert, R., & Krügel, A. (2010). Readers use Bayesian estimation for eye movement control. *Psychological Science*, 21(3), 366–371.
- Gigerenzer, G., & Brighton, H. (2009). Homo Heuristicus: Why biased minds make better inferences. *Topics in Cognitive Science*, 1, 107–143.
- Gigerenzer, G., Hoffrage, U., & Goldstein, D. G. (2008). Fast and frugal heuristics are plausible models of cognition: Reply to Dougherty, Franco-Watkins, and Thomas. *Psychological Review*, 115(1), 230–239.
- Kellerer, H., & Pferschy, U. (1999). A new fully polynomial time approximation scheme for the knapsack problem, *Journal of Combinatorial Optimization*, 3(1), 59–71.
- Koller, D., & Friedman, N. (2009). *Probabilistic graphical models: Principles and techniques*. Cambridge, MA: MIT Press.
- Kruschke, J.K. (2010). Bridging levels of analysis: comment on McClelland et al. and Griffiths et al. *Trends in Cognitive Sciences*, 14(8), 344–345.
- Kwisthout, J. (2009). *The computational complexity of probabilistic networks*. Doctoral dissertation, Department of Computer Science, Utrecht University, 2009.
- Kwisthout, J. (2010). Two new notions of abduction in Bayesian networks. *Proceedings of the 22nd Benelux Conference on Artificial Intelligence*, pp. 82–89.
- Kwisthout, J. (2011). Most probable explanations in Bayesian networks: Complexity and tractability. *International Journal of Approximate Reasoning*, 52, 1452–1469.
- Kwisthout, J. (2012). Structure approximation of most probable explanations in Bayesian networks. *Proceedings of the 24th Benelux Conference on Artificial Intelligence*.
- Kwisthout, J., & Van der Gaag, L.C. (2008). The computational complexity of sensitivity analysis and parameter tuning. *Proceedings of the 24th Conference on Uncertainty in Artificial Intelligence*, pp. 349–356.
- Kwisthout, J., Wareham, T., & van Rooij, I. (2011). Bayesian intractability is not an ailment that approximation can cure. *Cognitive Science*, 35(5), 779–784.
- Motwani, R., & Raghavan, P. (1995). *Randomized Algorithms*. New York: Cambridge University Press.
- Müller, M., van Rooij, I., & Wareham, T. (2009). Similarity as tractable transformation. *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, pp. 50–55.
- Oaksford, M., & Chater, N. (2009). Precise of Bayesian rationality: The probabilistic approach to human reasoning. *Behavioral And Brain Sciences*, 32, 69–120.
- Park, J.D. (2002). Using weighted MAX-SAT engines to solve MPE. *Proceedings of the Eighteenth National Conference on Artificial Intelligence*, pp. 682–687.
- Park, J.D., & Darwiche, A. (2004). Complexity results and approximation settings for MAP explanations. *Journal of Artificial Intelligence Research*, 21, 101–133.
- Sanborn, A.N., Griffiths, T.L., & Navarro, D.J. (2010). Rational approximations to rational models: Alternative algorithms for category learning. *Psychological Review*, 117(4), 1144–1167.
- Shimony, S.E. (1994). Finding MAPs for belief networks is NP-hard. *Artificial Intelligence*, 68(2), 399–410.

- Sloman, S.A., & Hagmayer, Y. (2006). The causal psycho-logic of choice. *Trends in Cognitive Sciences*, 10, 407–412.
- Stewart, N., Chater, N. & Brown, G.D.A. (2006). Decision by sampling. *Cognitive Psychology*, 53, 1–26.
- Valiant, L. (1984). A theory of the learnable. *Communications of the ACM*, 27(11), 1134–1142.
- van Rooij, I. (2008). The tractable cognition thesis. *Cognitive Science*, 32(6), 939–984.
- van Rooij, I., Evans, P., Müller, M., Gedge, J. & Wareham, T. (2008). Identifying sources of intractability in cognitive models: An illustration using analogical structure mapping. *Proceedings of the 30th Annual Conference of the Cognitive Science Society*, pp. 915–920.
- van Rooij, I., Kwisthout, J., Blokpoel, M., Szymanik, J., Wareham, T., & Toni, I. (2011). Intentional communication: Computationally easy or difficult? *Frontiers in Human Neuroscience*, 5(52), 1–18.
- van Rooij, I., & Wareham, T. (2008). Parameterized complexity in cognitive modeling: Foundations, applications and opportunities. *The Computer Journal*, 51, 385–404.
- van Rooij, I., & Wareham, T. (2012). Intractability and approximation of optimization theories of cognition. *Journal of Mathematical Psychology*, 56, 232–247.
- Vul, E., Goodman, N.D., Griffiths, T.L., & Tenenbaum, J.B. (2009). One and done? Optimal decisions from very few samples. *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, pp. 148–153.
- Wolpert, D.M., & Ghahramani, Z. (2005). Bayes rule in perception, action and cognition. In R.L. Gregory (Ed.), *Oxford companion to consciousness*. New York: Oxford University Press.
- Yuille, A., & Kersten, D. (2006). Vision as Bayesian inference: analysis by synthesis? *TRENDS in Cognitive Sciences*, 10(7), 301–308.
- Zuckerman, D. (2006). Linear degree extractors and the inapproximability of max clique and chromatic number, *Proceedings of the 38th ACM Symposium on Theory of Computing*, pp. 681–690.




Figure 1: (left) The most probable value assignment for V_1 to V_4 (denoted by the light- middle- and dark-grey colors); (middle) an approximation that resembles the most probable value assignment but may have a much lower probability; (right) an approximation that may have an almost-as-high probability as, but does not resemble, the most probable value assignment




Figure 2: A network with low (left) and high (right) treewidth. Note that there is only a single connection between the clusters A,B,C and D,E,F in the left picture, whereas the connections in the right picture are much more scattered over the whole network.