

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

Mark Blokpoel^{a,b}, Johan Kwisthout^a, Theo P. van der Weide^a, Iris van Rooij^b
(blokpoel@acm.org), (johank@science.ru.nl), (tvdw@cs.ru.nl), (i.vanrooij@donders.ru.nl)

^aRadboud University Nijmegen, Institute for Computing and Information Sciences

^bRadboud University Nijmegen, Donders Institute for Brain, Cognition and Behaviour

Abstract

An important aspect of human sociality is our ability to understand the actions of others as being goal-directed. Recently, the now classic rational approach to explaining this ability has been given a formal incarnation in the Bayesian Inverse Planning (BIP) model of Baker, Saxe, and Tenenbaum (2009). The BIP model enjoys considerable empirical support when tested on ‘toy domains’. Yet, like many Bayesian models of cognition, it faces the charge of computational intractability: i.e., the computations that the model postulates may be too resource demanding for the model to be scalable to domains of real-world complexity. In this paper, we investigate ways in which the BIP model can possibly parry the charge. We will show that there are specific conditions under which the computations postulated by the model are tractable, despite the model being rational and Bayesian.

Keywords: *goal inference, inverse planning, computational complexity, intractability, NP-hard, fixed-parameter tractability*

Introduction

Imagine a mother and her son, sitting in the same room, when she hears his stomach rumble. She sees her son get up, walk to the kitchen and start searching for something. At first he finds a sour apple, which he discards in search of something else. Then the mother sees her son finding a delicious candy bar. When he starts to eat it she realizes her son is trying to still his hunger and at the same time wanting to eat something sweet. In this scenario, the son goes through a process of *planning*, choosing his actions to achieve his goals. The mother observes the actions of her son and based on her observations infers the goals she thinks her son is trying to achieve. This process is called *goal inference*.

In line with a long tradition of explaining the human ability to understand actions as goal-oriented (Dennett, 1987; Charniak & Goldman, 1991; Csibra, Gergely, Biró, Koós, & Brockbank, 1999; Cuijpers, Schie, Koppen, Erlhagen, & Bekkering, 2006), Baker, Saxe, and Tenenbaum (2009) have proposed that goal inference can be seen as a form of *inverse planning*, just as vision is a form of inverse graphics. Baker et al. go beyond existing psychological approaches by providing a precise formalization of ‘inverse planning’ in the form of a Bayesian inference model. We will refer to this model as the BIP model of goal inference (where BIP stands for Bayesian Inverse Planning). The BIP model has been tested in several experiments, and Baker et al. (2007, 2009) observed that it can account for the dynamics of goal inferences made by human participants in several different experimental settings.

According to the BIP model, observers assume that actors are ‘rational’ in the sense that they tend to adopt those actions

that best achieve their goals. Given the assumption of rationality, and (probabilistic) knowledge of the world and how actions are affected by it, one can compute the probability that an agent performs an action given its goals, denoted

$$P(\text{action} \mid \text{goal}, \text{environment}) \quad (1)$$

When observing a given action, the probability in (1) can be inverted using Bayes’ rule to compute the probability of a given goal:

$$P(\text{goal} \mid \text{action}, \text{environment}) \propto$$

$$P(\text{action} \mid \text{goal}, \text{environment}) P(\text{goal} \mid \text{environment}) \quad (2)$$

Of all the possible goals that an observer can (or does) entertain, the goal that maximizes the probability in (2) best explains why the observed action was performed and is the goal that is inferred.¹

Given that the BIP model belongs to the class of (rational) Bayesian inference models—and Bayesian inference is known to be intractable if no additional constraints are imposed (e.g. Chater, Tenenbaum, and Yuille (2006); see also Kwisthout (2009))—the question arises if the computations that it postulates can scale to situations of everyday complexity. As Gigerenzer and colleagues put it:

The computations postulated by a model of cognition need to be tractable in the real world in which people live, not only in the small world of an experiment with only a few cues. This eliminates NP-hard models that lead to computational explosion, such as probabilistic inference using Bayesian belief networks ... including its approximations. (Gigerenzer, Hoffrage, and Goldstein (2008) p. 236)

Although we share the stance of Gigerenzer et al. (2008) towards intractable (NP-hard) models of cognition, we are not as pessimistic about the viability of Bayesian models. In our view, the key to understanding the computational feasibility of a Bayesian (or any cognitive) model lies in studying domain-specific constraints that hold in the model’s domain of application (e.g., action understanding or vision) and investigating if and how such constraints may render the computations postulated by the model tractable for its domain,

¹In other words, in the BIP model, goal inference is conceptualized as a form of probabilistic inference to the best explanation, a.k.a. *abduction* (e.g. Charniak and Shimony (1990)).

despite the intractability of those computations in general. In this paper we set out to perform such an investigation for the BIP model of goal inference.²

The remainder of this paper is organized as follows. We first introduce specific versions of the BIP model that Baker et al. (2007, 2009) formulated to account for their experimental data and observe that these versions are tractable but also too specific. We then propose a generalized model that breaks some implausible constraints in the original models. After this we introduce a method that we use to analyze the computational (in)tractability of the generalized model. We then give an overview of the (in)tractability results, and discuss their implications for Bayesian models of goal inference and for dealing with the intractability of Bayesian models in general.

Computational Models

Baker et al. (2009) propose three different versions of Bayesian Inverse Planning (M1, M2 and M3) to account for data gathered in several maze experiments. These two-dimensional maze experiments, based on earlier work (Gergely, Nádasdy, Csibra, & Biró, 1995; Schultz et al., 2003), were designed to assess subjects' inferences about the goals of a planning agent. Subjects were shown videos of agents moving in a maze, such as those in Fig. 1, and under different timing and information conditions had to infer the goal of the agent. In these experiments *changes in location* were considered *actions* and the *location* of the agent is considered its *state*. Specific locations (A, B and C) were possible *goals*.

All three models M1–3 can be seen as special cases of a more general BIP model, as depicted in Fig. 2, in which there is a goal structure template G that can encode different types of goal structures.³ The simplest goal structure is present in M1 where the observer assumes that the agent has one single goal that does not change over time (Fig. 3(a)). In M2 the model allows the observer to infer the agent has a different goal at any given time (Fig. 3(b)). This models the ability of people to infer changes in an agent's goal over time. For instance, if someone is inspecting the contents of her fridge, you may infer she wishes to cook dinner, but when she closes the fridge, puts on her coat, and leaves the house, you may

²The authors are well aware of common claims of approximability of Bayesian inferences, and that approximation is generally believed to provide a way to overcome the intractability of Bayesian models. In this paper, we will depart from this standard viewpoint for two reasons. First, the claims of approximability seem at worst incorrect and at best unfounded; for instance it is known that approximating the most probable explanation in a Bayesian network is itself also intractable (Abdelbar & Hedetniemi, 1998). Second, we believe that there are other, better ways of dealing with the intractability of cognitive models, viz., by identifying model constraints that render otherwise intractable models tractable (van Rooij, 2008).

³In the original BIP models (M1, M2 and M3) Baker et al. used additional parameters to model the effect of noise (β), the probability of changing a goal in M2 (γ) and the probability of having sub-goals in M3 (κ) to fit the model to the experimental data. As these parameters are assumed constants, they can be safely ignored for the purposes of our analyses.




Figure 1: An illustration of the types of stimuli used in the maze experiments of Baker et al. (2009). Participants observe an agent (and the trail history as memory aid) move inside the maze, and are asked to judge which of the three possible goals (A, B or C) is most likely the agent's goal. Here (a) depicts an early judgement point where both human participants and the model infer B as most likely goal. (b) depicts a later judgement point where both human participants and the model infer A as most likely goal. (c) A possible BIP model (in this case M1) for the early judgement point.

infer she is going to eat out. Finally, in M3 the goal structure encodes hierarchical goals (Fig. 3(c)), such that the observer can infer changes in the agent's sub-goals, which are subserving a common high-level goal. For instance, when you see someone gathering kitchen utensils, each individual gathering can be a sub-goal but the high-level goal is to cook dinner.

Even though inference in Bayesian networks is hard in general, the BIP models proposed by Baker et al. are tractable.⁴ This tractability is in some sense an artifact of the simplified experiments for which these models were designed. In the experiments an agent never has more than one (high-level) goal at any given time. This property does not seem to hold in general, however. Reconsider, for instance, the scenario in our opening paragraph. There the mother infers that the son wants to satisfy his hunger *and* he wants to eat something sweet. This type of goal inference where multiple goals are inferred at the same time cannot be modelled by M1, M2 or M3. To accommodate for this observation, we propose an extension called MULTIPLE GOALS BIP or MGBIP. Fig. 4

⁴For the formal proof of these claims we refer the reader to the Supplementary materials available online at <http://tinyurl.com/suppl2010>




Figure 2: A graphical representation of the dynamic Bayesian network that describes the general form of BIP. Nodes represent variables, for example state node S_t is a variable and it can assume values corresponding to the state of the agent at time t . Arrows represent dependencies, for example the probability that a state S_{t+1} has a certain value depends on the previous state S_t and previous action A_t . States and actions are observed, i.e. the values of the states and action variables are given as input to the model. Given these observations the most probable combination of values for the goal variables in \mathbf{G} . Finally, shaded nodes are observed and their values considered part of the input of the model. Examples of the possible contents of \mathbf{G} are illustrated in Fig. 3


illustrates the dynamic Bayesian network of MGBIP.⁵

Because it is more general, MGBIP has wider range of applicability than M1–3. The introduced generality also comes at a cost: Whereas M1, M2 and M3 are tractable, MGBIP is intractable, in the sense that there are no tractable (more precisely: polynomial time) algorithms that can implement this model.⁴ Even so, in real-world situations humans are often able to quickly infer an agent is pursuing multiple simultaneous goals. This suggests that, if MGBIP is to be psychologically plausible, we need to assume that some domain-specific constraints apply in those situations that render the goal inferences tractable under the MGBIP model (despite the model being intractable without such additional constraints). The next section describes how we set out to identify such possible constraints.


Identifying Sources of Intractability

In order to find constraints on the input domain of MGBIP that render the (restricted) model tractable, we adopt a method for identifying sources of intractability as described in (van Rooij, Evans, Müller, Gedge, & Wareham, 2008) (see also van Rooij and Wareham (2008)). The method works as follows.


⁵Baker et al. (2009) also note, that the simplified models M1–3 unlikely suffice to model human action understanding in general and they argue that the models will need to be extended in various directions if they are to apply to real-world scenarios. Our extension can be seen as one such possible direction in which to extend the model. Other directions of extension are possible as well (see e.g. (Ullman, Baker, Macindoe, Goodman, & Tenenbaum, 2009).



(a) M1



(b) M2



(c) M3

Figure 3: Graphical representation of \mathbf{G} for M1, M2 and M3. In M1 (a) goals are modeled by a single static goal. All actions are dependent on this goal. In M2 (b) goals can change over time. Actions at time t are dependent on goals at time t . In M3 (c) goals can consist of multiple subgoals. Actions at time t are dependent on subgoals at time t .

First, one identifies a set of model parameters $K = \{k_1, k_2, \dots, k_m\}$ in the model M under study (for us, MGBIP). Then one tests if it is possible to solve M in a time that can grow excessively fast (more precisely: exponential or worse) as a function of the elements in K yet slowly (polynomial) in the size of the input.⁶ If this is the case, then M is said to be *fixed-parameter (fp-) tractable* for parameter set K , and otherwise it is said to be *fp-intractable* for K .

Observe that if a parameter set K is found for which M is fp-tractable then the problem M can be solved quite efficiently, even for large inputs, provided only that the members

⁶More formally, this would be a time on the order of $f(k_1, k_2, \dots, k_m)n^c$, where f is an arbitrary computable function, n is a measure of the overall input size, and c is a constant.




Figure 4: Graphical representation of the dynamic Bayesian network that describes MULTIPLE GOALS BIP (MGBIP).

of K are relatively small. In this sense the “unbounded” nature of K can be seen as a reason for the intractability of M . Therefore we call K a *source of intractability* of M .

The MGBIP model has several natural parameters, each of them a candidate source of intractability. In this paper we consider, five such parameters (see Table 1 for an overview and Fig. 5 for an illustration).

First consider parameters T , denoting the maximum number of observations the observer makes, and $1/T$, denoting the poverty of observations. Note that T is small if few observations are made, and $1/T$ is small if many observations are made. Based on intuition one might think, the less information we have, the harder it is to understand actions. This makes $1/T$ a candidate source of intractability. However as T grows, so does the size of the network and the necessary number of calculations and this also makes T a likely candidate source of intractability.

Second, parameter k is the maximum number of multiple goals that (the observer assumes) the agent can pursue. This parameter is also an excellent candidate source of intractability, because large k 's introduce an exponential number of combinations of possible multiple goals leading to a combinatorial explosion.

Third, the parameter g is the maximum number of goal values per goal variable. As the number of possible values that a goal variable can take increases the necessary number of calculations, also g is a candidate source of intractability.

Finally, the parameter $1 - p$ measures how far the most likely goal inference is from being completely certain (here p is the probability of the most likely explanation). If $1 - p$ is small, this means that the most likely explanation is much more likely than any competitor explanation. If the value is large, it means that the most likely explanation has many competitor explanations of non-negligible probability (see e.g. Table 2). It seems intuitive that finding the most likely explanation is easier in the former case than in the latter case, and therefore also $1 - p$ can be considered a candidate source of intractability.

Table 1: A list of parameters with short descriptions and their values based on the running example.

<i>parameter</i>	<i>description</i>	<i>value</i>
T	maximum observations	6
$1/T$	maximum observation poverty	1/6
k	maximum # multiple goals	2
g	maximum # goal values	3
$1 - p$	distance from certainty	0.4

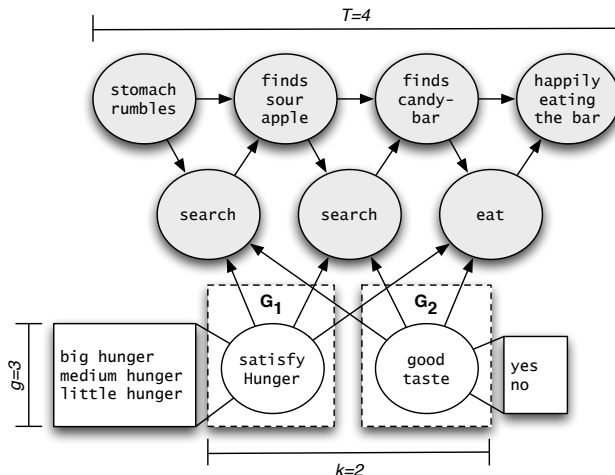


Figure 5: Illustration of the Bayesian network and different parameters of the MGBIP model applied to the “mother observes son”-example.

Table 2: Example probability distribution over the combinations of goal values. In this example $p = 0.6$ and $1 - p = 0.4$.

	satisfy hunger	desire sweet	P
big hunger	yes	0.05	0.05
medium hunger	yes	0.05	0.05
little hunger	yes	0.6	0.6
big hunger	no	0.3	0.3
medium hunger	no	0.0	0.0
little hunger	no	0.0	0.0

We have now reviewed five parameters that—on intuitive grounds—may be considered candidate sources of intractability in the MGBIP model. It is known, however, that human intuitions about what makes a computation tractable or intractable can be mistaken. Therefore it is necessary to verify such intuitions by means of mathematical proof.

Results

In this section we present our fp-(in)tractability results for the different parameters of the MGBIP model, and we explain how these results bear on the question ‘which constraints render the MGBIP model tractable?’. Full details and proofs can

be found in the Supplementary materials.⁷

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

Result 1 shows—contrary to the intuitions sketched in the previous section—that none of the parameters T , $1/T$ and g , nor any combination of them is a source of intractability for MGBIP. This means that even if we assume that one or more of these parameters is small for the domain of application, goal inference under the MGBIP model is still intractable.

Besides this negative result (Result 1), we also have two positive results (Results 2 and 3).

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

Result 2 confirms parameter k is a source of intractability. This means that goal inference is tractable under the MGBIP model provided only that we impose the constraint that (the observer assumes that) the agent can pursue only a handful of goals simultaneously. Importantly, this is true regardless the size of T , $1/T$, g or $1 - p$. This is quite a powerful result, with great potential for explaining the speed of real-world goal inferences within the confines of a BIP model. After all, it seems to be a plausible constraint that real-world agents do not (typically) pursue a large number of goals in parallel at the same time (possibly also to keep their own planning tractable).

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

Finally, Result 3 confirms parameter $1 - p$ is a source of intractability. This means that goal inference is tractable under the MGBIP model for those inputs where the most probable goal explanation is quite probable. Again, this is true regardless the size of T , $1/T$, g or k . Also, this result has potential for explaining the speed of real-world goal inferences within the confines of a BIP model, at least for certain situations—viz., those situations where the actions of the observed agents unambiguously suggest a particular combination of goals. For all we know, real world cases of speedy goal inference may very well match exactly these situations. Whether or not this is indeed the case is an empirical question which can be addressed by testing the speed of human goal inference for different degrees of goal ambiguity.

Discussion

We have analyzed the computational resource requirements of the Inverse Bayesian Planning (BIP) model of goal inference in order to study its viability as a model of inferences made by resource-bounded minds as our own. We generated several interesting theoretical findings. First, we observed that the three specialized models—M1, M2, and M3—that were developed by Baker et al. (2007, 2009) to account for their experimental data in maze experiments are in fact computationally tractable. This means that these specialized

Bayesian models do not seem to make unrealistic assumptions about the computational powers of human minds/brains, even when operating on large networks of beliefs and observations. That being said, these models do seem to be theoretically problematic for a different reason: they are too specialized to count as models of goal inference in general.

The over-specialization of M1, M2 and M3 is revealed when pondering the assumptions that these models make about the agent and the observer. For instance, all three models assume that (the observer assumes that) the agent can pursue at most one goal at a time. In the real-world, however, people often can and do act in ways so as to try and achieve two or more goals at the same time, and observers can also often understand what these simultaneous goals are from observing the actors behave in systematic ways. Recall, for example, the scenario from our Introduction where the son searches the kitchen for a candy bar. Under different circumstances, the mother may understand that her son has the goal to still his hunger (goal 1), to satisfy his craving for sweet (goal 2), to see how many bars are left (goal 3), to pretend that he did not hear his mom's request to clean up his room (goal 4), to bring back a candy bar for his mom (goal 5), etc., or any combination of these goals.

To accommodate the fact that real-world goal inference is not restricted to one goal at a time, we defined a more general BIP model—having M1, M2 and M3 as special cases—which we refer to as MULTIPLE GOALS BIP, or MGBIP for short. Complexity Analysis of the MGBIP model revealed that it *is* computationally intractable (i.e., NP-hard), meaning that this model, in all its generality, does indeed make unrealistic assumptions about the computational powers of human minds/brains. We took this negative theoretical result to mean that—if the BIP model is to account for human goal inference at all—it must be the case that in those situations where humans are able to infer multiple simultaneous goals quickly and effortlessly, specific constraints apply that render the inferences under the MGBIP model tractable.

To investigate which types of constraints could render the MGBIP model tractable, we used a methodology for identifying sources of intractability in NP-hard computational models (e.g. van Rooij and Wareham (2008)) and derived several theoretical results. For instance, we ruled out the possibility of explaining speedy real-world (multiple) goal inferences by an appeal to small values of T (modeling situations when goals can be inferred using only few observations) or an appeal to large values of T (modelling situations where a lot of information is available on which to base a goal inference). Similarly, we ruled out that the speed of such inferences could be explained by an appeal to a small number of values per goal node. Besides these negative theoretical results, we also had two important positive results. For one, we established that as long as the number of goals that can be simultaneously pursued, k , is not too large then goal inference is tractable under the MGBIP. Secondly we have shown that goal inference is tractable under the MGBIP model whenever the probability

⁷See <http://tinyurl.com/suppl2010>

of the most likely combination of simultaneous goals, p , is not too far from 1.

Whereas our negative theoretical results are useful to clarify that tractability is not a property that is trivially achieved (and often our intuitions about what constraints would render a model tractable can be wrong; cf. van Rooij et al. (2008)), our positive results show that a model of action understanding can nevertheless be rational, Bayesian, and tractable. Moreover, the nature of the constraints that need to be introduced to render the Bayesian Inverse Planning model of goal inference tractable yield new empirically testable predictions.

For instance, based on our results, we predict that human participants will be able to make quick and accurate goal inferences in the types of experimental set-ups such as used by Baker et al. (2007) (but see also Csibra et al. (1999)), but only if the number of simultaneous goals that the observed agents are pursuing is not too large, or the probability of the most likely combination of goals is not too small, or both. If both of these constraints were to be alleviated at the same time, we would predict that human performance on the goal inference task would deteriorate significantly. If our prediction were to be confirmed then this would provide corroborative support for the BIP model of goal inference, and validate that our theoretical results help explain the tractability of human goal inferences. If, on the other hand, the prediction were to be disconfirmed, then this would suggest that either the BIP model fails as an account of human goal inferences, or some constraint other than the ones we considered also suffices to render the BIP model tractable. The latter option may then be one that BIP modelers may be interested in pursuing further.

In closing, we remark that our approach can be seen as exemplary of a general strategy for dealing with intractability in Bayesian models, whether of action understanding or otherwise. Our approach reveals that—contrary to popular belief—Bayesian models *can* possibly scale to complex, real-world domains. To achieve this, Bayesian modelers need only identify constraints that apply in the real-world and suffice to render their models' computations tractable. By restricting Bayesian models in this way these models also become better testable: the constraints required to guarantee tractability of the models yield new predictions (specifically, about the speed of inferences) that can be used to perform more stringent tests of such models.

Acknowledgments

The 2nd author has been supported by the OCTOPUS project under the responsibility of the Embedded Systems Institute. The OCTOPUS project is partially supported by the Netherlands Ministry of Economic Affairs under the Embedded Systems Institute program.

The authors thank Vanessa Ferdinand and Max Hinne for comments on an earlier version of this paper.

References

Abdelbar, A. M., & Hedetniemi, S. M. (1998). Approximating MAPs for belief networks is NP-hard and other theo-

- rems. *Artificial Intelligence*, 102, 21–38.
- Baker, C. L., Saxe, R., & Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition*(113), 329–349.
- Baker, C. L., Tenenbaum, J., & Saxe, R. (2007, Jan). Goal inference as inverse planning. *Proceedings of the 29th meeting of the Cognitive Science Society*.
- Charniak, E., & Goldman, R. P. (1991). A probabilistic model of plan recognition. In *Association for the advancement of artificial intelligence* (pp. 160–165).
- Charniak, E., & Shimony, S. E. (1990). Probabilistic semantics for cost based abduction. In *Aaai* (p. 106–111).
- Chater, N., Tenenbaum, J. B., & Yuille, A. (2006). Probabilistic models of cognition: Conceptual foundations. *Trends in Cognitive Sciences*, 10(7), 287 – 291. (Special issue: Probabilistic models of cognition)
- Csibra, G., Gergely, G., Biró, S., Koós, O., & Brockbank, M. (1999). Goal attribution without agency cues: the perception of 'pure reason' in infancy. *Cognition*, 72, 237–267.
- Cuijpers, R. H., Schie, H. T. van, Koppen, M., Erlhagen, W., & Bekkering, H. (2006). *Goals and means in action observation: A computational approach*.
- Dennett, D. C. (1987). *The intentional stance*. Cambridge, MA: The MIT Press.
- Gergely, G., Nádasdy, Z., Csibra, G., & Biró, S. (1995). Taking the intentional stance at 12 months of age. *Cognition*, 56, 165–193.
- Gigerenzer, G., Hoffrage, U., & Goldstein, D. G. (2008). Fast and frugal heuristics are plausible models of cognition: Reply to Dougherty, Franco-Watkins, and Thomas (2008). *Psychological Review*, 115, 230–239.
- Kwisthout, J. H. P. (2009). *The computational complexity of probabilistic networks*. Unpublished doctoral dissertation, Faculty of Science, Utrecht University, The Netherlands.
- Schultz, R. T., Grelotti, D. J., Klink, A., Kleinman, J., Gaag, C. van der, & Marois, R. (2003). The role of the fusiform face area in social cognition: Implications for the pathobiology of autism. *Philosophical Transactions of Royal Society of London, Series B: Biological Sciences*(358(1430)), 415–427.
- Ullman, T. D., Baker, C. L., Macindoe, O., Goodman, O. E. N. D., & Tenenbaum, J. B. (2009). Help or hinder: Bayesian models of social goal inference. *NIPS*.
- van Rooij, I. (2008). The tractable cognition thesis. *Cognitive Science*, 32, 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. In K. M. B. C Love & V. M. S. (Eds.) (Eds.), *Proceedings of the 30th annual conference of the cognitive science society* (pp. 915–920). Austin, TX: Cognitive Science Society.
- van Rooij, I., & Wareham, T. (2008). Parameterized complexity in cognitive modeling: Foundations, applications and opportunities. *Computer Journal*, 51(3), 385–404.