



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

Mark Blokpoel

Submitted in total fulfilment of the requirements  
of the degree of Master of Science

May 14, 2010

Radboud University

Theo P. van der Weide<sup>a</sup>, Iris van Rooij<sup>b</sup>

<sup>a</sup>Radboud University, Institute for Computing and Information Sciences

<sup>b</sup>Radboud University, Donders Institute for Brain, Cognition and Behaviour



## Cover

We observe actions of other humans all the time and we are able to figure out those peoples goals quickly. But when you think about it, it is not that easy. Consider the actions on the cover from upper-right to bottom-left:

1. Will he shoot the ball left or right?
2. Is she feeding or catching the duck?
3. Is she stretching or dancing?
4. Is she resting or waiting for someone?
5. Is he warning us or happy he won the race?
6. Does he want to correct her or is he flirting?
7. Are they dancing or fighting?
8. Is he helping the child walk or is he punishing him?
9. Is he trying to clean the dishes or bring drinks?

1, 2, 4, 5, 6, 7, 8 and 9 by Michal Zacharzewski  
3 by Kymberly Vohsen



# Preface

This thesis marks the end of my study in Computer Science and the beginning of a new study in Cognitive Science. About one year ago I met Iris, when I participated in a course she teaches called Cognition & Complexity. I learned to use methods from computer science to explain the speed of human cognitive capabilities. “How is it possible we can quickly understand the analogy between a composer and a general?” “Why can we quickly assess similarity between objects?” Each time I understood the answer to these kinds of questions I had a eureka moment and oh, how I loved those moments. . .

After the course I was hooked and I was left with the question: “How is it possible that humans can quickly understand each other?” It looks as if there are more explanations for observed actions than there are stars in the universe and finding the best solution seems a daunting – if not impossible – task. Yet, humans can do it in split seconds! The thesis in front of you describes the first steps towards an answer to this question.

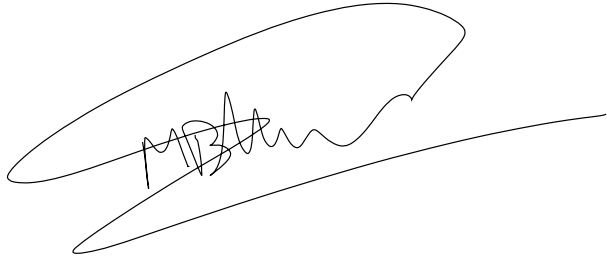
My thesis could not have been possible without the help and collaboration of my colleagues. A big thanks goes to Johan for his hard work and extremely quick e-mail responses. I also wish to thank my fellow interns at the Donders Institute, for their interest in me and my research. A special thanks goes to Jop and Max: your comments, willingness to listen and motivational insights were very helpful and dear to me.

Next I wish to thank the people who helped shape me as a scientist. Franc, for nurturing my curiosity and spirit throughout my study. You saw a scientist in me years before I realized it myself. Theo, for his supervision: your guidance not only got the best out of me but also helped me decide what career to choose. And last but certainly not least, Iris. Your unparalleled enthusiasm, interest and pride in both our research and me personally were really contagious and made me feel part of the community.

Finally I wish to thank my friends and family for believing in me and their unconditional support. Mom, dad, Erwin: You never gave me the feeling I needed to make you feel proud, because you are proud of me no matter what. That means a lot to me and makes me feel proud of what I accomplished and of being your son and brother. Tom, Frank: You guys are

---

always interested even when the conversation got theoretical. It is always a pleasure to discuss my research with you. Most importantly I wish to thank Wietske. Your care, love and support make me feel special. You listen to my ramblings even though they don't make sense. To see you proud and smile is the biggest reward of all.

A handwritten signature in black ink, appearing to be 'M. B.', is written over a large, loopy, horizontal scribble that spans across the middle of the page.

# Contents

<b>1</b>	<b>Introduction</b>	<b>9</b>
<b>2</b>	<b>Preliminaries</b>	<b>13</b>
2.1	Bayesian modeling . . . . .	13
2.2	Computational complexity . . . . .	15
2.2.1	Traditional computational complexity . . . . .	15
2.2.2	Parameterized complexity . . . . .	17
<b>3</b>	<b>Computational Models</b>	<b>19</b>
3.1	M1, M2 and M3 are tractable . . . . .	21
3.2	MULTIPLE GOALS BIP . . . . .	23
3.3	MGBIP is intractable . . . . .	24
<b>4</b>	<b>Identifying Sources of Intractability</b>	<b>31</b>
4.1	MGBIP fp-intractability . . . . .	33
4.2	MGBIP fp-tractability . . . . .	35
<b>5</b>	<b>Discussion</b>	<b>39</b>
<b>6</b>	<b>Bibliography</b>	<b>43</b>
	References . . . . .	43





# 1

## 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 wants 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 (Baker, Tenenbaum, & Saxe, 2007; Baldwin & Baird, 2001; Cuijpers, Schie, Koppen, Erlhagen, & Bekkering, 2006; Hassin, Aarts, & Ferguson, 2005; Király, Jovanovic, & Prinz, 2003; van Rooij & Wareham, 2008), Baker, Saxe, and Tenenbaum (2009) have proposed that goal inference can be seen as a form of *inverse planning*, just as vision is believed to be 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 ac-

count 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 effected by it, one can compute the probability that an agent performs an action given its goals, denoted

$$\Pr(\textit{action} \mid \textit{goal}, \textit{environment}) \tag{1.1}$$

When observing a given action, the probability in Equation 1.1 can be inverted using Bayes’ rule to compute the probability of a given goal:

$$\Pr(\textit{goal} \mid \textit{action}, \textit{environment}) \propto \Pr(\textit{action} \mid \textit{goal}, \textit{environment}) \Pr(\textit{goal} \mid \textit{environment}) \tag{1.2}$$

Of all the possible goals that an observer can (or does) entertain, the goal that maximizes the probability in Equation 1.2 best explains why the observed action was performed and is the goal that is inferred. In other words, in the BIP model, goal inference is conceptualized as a form of probabilistic inference to the best explanation, also known as *abduction* (e.g. Charniak and Shimony (1990)).

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 J. H. P. 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, despite the intractability of those models in general. In this thesis we set out to perform such an investigation for the BIP model of goal inference.

The methodology we use allows us to identify domain-specific constraints that render otherwise intractable models tractable (van Rooij & Wareham, 2008). We depart from the standard view that approximability of Bayesian inferences (or other models) can overcome the intractability of models 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, when tractability of a model is claimed through approximation, then a more accurate model of the cognitive process is the approximation model, not the original model. The approximation model should then be explicated and is still subject to the tractability issue.

The remainder of this thesis is organized as follows. First, Chapter 2 are preliminaries to introduce the topics of Bayesian models and Computational complexity theory. Second, in Chapter 3 we introduce specific versions of the BIP model Baker et al. (2007, 2009) formulated to account for their experimental data and observe that these versions are tractable but also too specific. We also propose a generalized model that breaks an implausible constraint in the original models. After this, in Chapter 4, we introduce a method that allows us to analyze the computational (in-)tractability of the generalized BIP model, we use this method to analyze the model and present the (in-)tractability results. Finally, in Chapter 5, we discuss their implications for Bayesian models of goal inference and for dealing with the intractability of Bayesian models in general.



# 2

## Preliminaries

In this chapter we review basic concepts from Bayesian modeling, computational complexity theory and the inverse planning framework. Readers unfamiliar with these concepts are advised to study this chapter as they are necessary for a good understanding of the following chapters.

### 2.1 Bayesian modeling

For readers unfamiliar with basic notations from Bayesian modeling we review some of the basics relevant for our purpose. In our notation capital letters ( $A, B, C, \dots$ ) denote variables, small letters ( $a, b, c, \dots$ ) denote values, bold letters ( $\mathbf{A}, \mathbf{a}, \mathbf{B}, \mathbf{b}, \dots$ ) denote sets and normal letters ( $A, a, B, b, \dots$ ) denote singletons. 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  $\Gamma = \{\Pr_X | X \in \mathbf{V}\}$  is the set of conditional probability distributions  $\Pr(X | \mathbf{y})$  for each joint value assignment  $\mathbf{y}$  to the parents of  $X \in \mathbf{G}$ . For clarity a BN is usually depicted by a graph, where directed edges  $(X, Y) \in \mathbf{A}$  represent dependencies  $\Pr(Y | X) \neq \Pr(Y)$ .

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  $W \in \mathbf{W}$  and

each associated value  $w \in \mathbf{w}$  such that  $\Pr(W = w) = 1$  and  $\Pr(W \neq w) = 0$ . 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  $\xi$  denote a Boolean algebra of propositions spanned by  $\mathbf{V}$ . The function  $\Pr : \xi \rightarrow [0, 1]$  is a joint probability distribution on  $\mathbf{V}$  if the following conditions hold:

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

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  $\forall_{t,i \in I} \exists_{o \in O} [\Pr(i_{t+1} \mid o_t) \in \Gamma]$ .

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 (as defined in Chapter 3) 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  $\mathbf{e}$ ?

Finally a *tree-decomposition* of a graph  $G = (V, E)$  is based on a set of tree-nodes (called bags)  $\mathcal{X} \subseteq \mathcal{P}(V)$  and a set of tree-edges  $F \subseteq \mathcal{X} \times \mathcal{X}$ , such that:

1.  $\mathcal{X}$  is a cover of  $V$ ,  $\cup \mathcal{X} = V$ ;
2. each edge in  $E$  is part of a set in  $\mathcal{X}$ ,  $\forall_{(x,y) \in E} \exists X \in \mathcal{X} [v \in X \wedge w \in X]$ ;
3. and each bag on a path between two bags contains the disjunction of those two bags,  $\forall_{(X,Y) \in F^+ \wedge (Y,Z) \in F^+} [X \cup Z \subseteq Y]$ .

The *treewidth* (Robertson & Seymour, 1986) of a BN  $\mathcal{B}$  is defined as the minimum width over all tree-decompositions of the moralized graph of  $\mathcal{B}$ , where the width of a tree-decomposition  $(\mathcal{X}, F)$  is equal to the size of a largest bag in  $\mathcal{X}$  minus 1,  $tw(\mathcal{B}) = \max_{X \in \mathcal{X}} |X| - 1$ .

## 2.2 Computational complexity

In the following chapters we also assume the reading is familiar with basic notions from computational complexity theory – this includes concepts such as Big-Oh  $O(\cdot)$ , (in-)tractability, polynomial time reductions and *NP*-hardness – and parameterized complexity theory – including concepts such as *fp*-(in-)tractability and parameter. This section is a short introduction to complexity theory, such that readers unfamiliar with the theory are able to read the remainder of the thesis. For full details on the theories we refer to textbooks (Garey & Johnson, 1979; Downey & Fellows, 1998).

Cognitive scientists try to model as best as possible an existing system capacity (namely the observed cognitive phenomenon). They often do this at the computational level (see (Marr, 1982)) by specifying the relation between the input and output domain of the phenomenon. In this thesis we use computational model – a term from cognitive psychology – and problem – a term from computer science – to denote the same concept: a function of some input to some output  $\Pi : I \rightarrow O$ .

### 2.2.1 Traditional computational complexity

In a computational complexity analysis we study the amount of computational resources – in our case time – required to compute the output of a problem  $\Pi$ . We express the complexity of a problem in terms of the

required resources as a function of the size of the input. We first define the Big-Oh notation used to express complexity. Big-Oh is an asymptotic upper-bound and we say a function  $f(x)$  is  $O(g(x))$ , if there are constants  $c \geq 0$  and  $x_0 \geq 1$  such that  $f(x) \leq cg(x)$  for all  $x \geq x_0$ . The Big-Oh notation ignores constants and low-order polynomials which is why it is also called the *order of magnitude*. For example  $x^3 + x^2 + x + 4$  is on the order of  $x^3$  or  $O(x^3)$  and  $1 + 2 + \dots + x = \frac{x(x+1)}{2}$  is  $O(x^2)$ .

We are interested in the time complexity of models in terms of the size of the input. The input  $i$  of a problem, model or function has size  $n = |i|$  which is the number of symbols used in a typical encoding (usually the input tape of a Turing Machine). A problem  $\Pi$  can be solved in time  $O(g(n))$  if there exists an algorithm that solves  $\Pi$  in time  $O(g(n))$ . The time complexity of  $\Pi$  is measured by the fastest *known* algorithm that solves  $\Pi$ .

Problems can be classified according to their nature and complexity into complexity classes such as  $P$  and  $NP$ . The class  $P$  contains all decision problems – problems that output only yes or no – that are solvable in polynomial time. A problem is solvable in polynomial time if there exists an algorithm that solves it in  $O(n^\alpha)$  for some constant  $\alpha$ . Class  $NP$  contains all problems that can be verified in polynomial time. Trivially  $P \subseteq NP$  and it is generally believed that  $P \neq NP$  (Sipser, 1992). A problem is *hard* for a certain complexity class  $C$  if it is at least as hard as all other problems in  $C$ . For example a problem  $\Pi$  is  $NP$ -hard if all other problems in  $NP$  are at least as hard as  $\Pi$ .

A problem  $\Pi$  is at least as hard as  $\Theta$  if there exists a polynomial time reduction from  $\Theta$  to  $\Pi$ . We say  $\Pi$  reduces to  $\Theta$  if there exists a function  $\tau$  that transforms any input  $i_\Pi$  of  $\Pi$  to input  $\tau(i_\Pi)$  of  $\Theta$  such that  $i_\Pi$  is a yes-instance for  $\Pi$  if and only if  $\tau(i_\Pi)$  is a yes-instance for  $\Theta$ . A reduction is a polynomial time reduction if  $\tau$  is polynomial time computable. We write  $\Pi \leq_\tau \Theta$  if  $\Theta$  polynomial time reduces to  $\Pi$ , i.e. if  $\Pi$  is at least as hard as  $\Theta$ . Polynomial time reductions are very powerful and can be used to prove problem is  $NP$ -hard or in  $P$ . If a problem  $\Theta$  is known  $NP$ -hard then  $\Pi$  is  $NP$ -hard if  $\Pi \leq_\tau \Theta$ . Vice-versa, if a problem  $\Pi$  is in  $P$ , then  $\Theta$  is also in  $P$  if and only if  $\Pi \leq_\tau \Theta$ .



### 2.2.2 Parameterized complexity

While traditional complexity theory provides a methodology to formalize the amount of required resources to solve a problem, it fails to detail what makes a problem (in-)tractable. In the 90s Downey and Fellows developed a variant on complexity theory called parameterized complexity theory (Downey & Fellows, 1998). Their framework expresses the complexity of a problem  $\Pi$  in terms of sets of parameters (or properties)  $\kappa$  of the input. If some set of these parameters has an exponential (or worse) contribution to the complexity of the problem, then the problem tractable if we assume the parameters in that particular set are upper-bounded by small values.

Let  $\Pi : I \rightarrow O$  be a problem,  $K$  be the set of all parameters of the input  $I$  and  $\kappa \subseteq K$ . We say  $\kappa$ - $\Pi$  is *fixed parameter tractable (fp-tractable)* if there exists at least one algorithm that computes  $O$  for all  $I$  in  $O(f(\kappa)n^\alpha)$ , where  $f$  is an arbitrary function of order exponential (or worse) and  $\alpha$  is a constant. If no such algorithm exists then  $\kappa$ - $\Pi$  is said to be *fixed parameter intractable (fp-intractable)*. Alternatively when  $\kappa$ - $\Pi$  is fp-(in)tractable we can say  $\Pi$  is fp-(in)tractable for  $\kappa$ .

Observe that if a parameter set  $\kappa$  is found for which  $\Pi$  is fp-tractable then the problem  $\Pi$  can be solved quite efficiently, even for large inputs, provided only that the members of  $\kappa$  are relatively small. In this sense the “unbounded” nature of parameters in  $\kappa$  can be seen as a reason for the intractability of  $\Pi$ . Therefore we call  $\kappa$  a *source of intractability* of  $\Pi$ .

The following lemmas in parameterized complexity are used in the proofs in this thesis.

**Lemma 2.1.** Let  $\Pi$  be a sub-problem of  $\Theta$ , where both  $\Pi$  and  $\Theta$  can be parameterized by  $\kappa$ . Then if  $\kappa$ - $\Theta$  is fp-tractable,  $\kappa$ - $\Pi$  is also fp-tractable.

**Lemma 2.2.** If a problem  $\Pi$  is fp-intractable for a parameter set  $\kappa$ , than  $\Pi$  is fp-intractable for any subset  $\kappa' \subseteq \kappa$ .

Confusion exists over related parameters such as  $a$  and  $1/a$ . Both parameters require separate tractability proofs as explained by the following lemma.

**Lemma 2.3.** Let  $\{a\}$ - $\Pi$  be computable in time  $O(f(a)n^\alpha)$ , where  $n$  is the size of the input,  $\alpha$  a constant and  $f(a)$  an exponential (or worse) growing function as  $a$  grows. Thus if  $a$  is upper-bounded,  $\Pi$  is tractable. Now let  $\{1/a\}$ - $\Pi$  be computable in time  $O(f(1/a)n^\alpha)$ , then  $f(1/a)$  is decaying as  $a$  grows and thus  $O(f(1/a)n^\alpha)$  cannot upper-bound the complexity of  $\{1/a\}$ - $\Pi$ . We need a function  $g(1/a)$  that grows as  $1/a$  grows to express the complexity of  $\{1/a\}$ - $\Pi$  as  $O(g(1/a)n^\alpha)$ .

# 3

## 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. 3.1, 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*. Figure 3.1(c) illustrates an example BIP model where NE and E are actions of stepping in that particular cardinal direction and  $(x, y)$  represent the location of the agent in the maze.

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 dynamic BN  $\mathcal{D}$  where each slice consists of a state variable  $S_t \in \mathbf{S}$  and action variable  $A_t \in \mathbf{A}$ . Additionally there is a set  $\mathbf{G}$  that contains an arbitrary number of variables that encode the goal(s). In this framework  $A_t$  depends on  $S_t$  and on (at least one) goal variable in  $\mathbf{G}$ . State variables  $S_{t+1}$  depend on the previous state  $S_t$  and action variable  $A_t$ . This means that for  $\mathcal{D}$ ,  $I_t = S_t$  and  $O_t = \{S_t, A_t\}$ .

In the original BIP models (M1, M2 and M3) Baker et al. used addi-




Figure 3.1: 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 judgment point where both human participants and the model infer A as most likely goal. (c) A possible BIP model for the early judgement point.

tional parameters to model the effect of noise ( $\beta$ ), prior probabilities based on world knowledge ( $w$ ), the probability of changing a goal in M2 ( $\gamma$ ) and the probability of having sub-goals in M3 ( $\kappa$ ) 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.

All three models M1–3 can be seen as special cases of a more general BIP model, as depicted in Fig. 3.2, in which there is a goal structure template  $G$  that can encode different types of goal structures. 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. 3.3(a)). In M2 the model allows the observer to infer the agent has a different goal at any given time (Fig. 3.3(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 infer




Figure 3.2: A graphical representation of the dynamic Bayesian network that describes the general form of BIP. States and actions are observed (depicted as shaded nodes), 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  $\mathbf{G}$  must be inferred. Examples of the possible contents of  $\mathbf{G}$  are illustrated in Fig. 3.3

she is going to eat out. Finally, in M3 the goal structure encodes hierarchical goals (Fig. 3.3(c)), such that the observer can infer changes in the agent's sub-goals, which are sub-serving a common high-level goal. For instance, when you see someone gather kitchen utensils, picking up a bowl and finding a spoon can be seen as sub-goals but the high-level goal is to cook dinner.

### 3.1 M1, M2 and M3 are tractable

Even though inference in Bayesian networks is hard in general, the BIP models proposed by Baker et al. are tractable. To prove M1–3 are tractable, we first define them as input/output-problems. In this definition we assume the model's output is the most likely joint value assignment to the variables in  $\mathbf{G}$ . Under this assumption M1, M2 and M3 are special cases – namely cases with restricted topology – of MPE. Figure 3.3 contains graphical




Figure 3.3: 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  $t$ .


representations of M1, M2 and M3.

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  $G_1, \dots, G_{T-1}$  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  $G_1, \dots, G_{T-1}$  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 J. Kwisthout (2010) for an overview) that solve MPE in polynomial time when the treewidth of the moralized graph of  $\mathcal{B}$  is small (i.e. they proved MPE is fp-tractable for treewidth). More in particular, the running-time is  $O(f(tw)g(p))$ , where  $f$  is an exponential function based



(a) M1 tree-decomposition (b) M2 tree-decomposition (c) M3 tree-decomposition

Figure 3.4: The tree-decompositions of M1–3.

on the treewidth of  $\mathcal{B}$  ( $tw$ ) and  $g$  is a polynomial based on the number of cliques in  $\mathcal{B}$  ( $p \leq |\mathbf{V}|$ ). Furthermore the following results of treewidth of M1–3 are known, based on the tree-decompositions in Figure 3.4.

<i>BIP model</i>	<i>treewidth</i>
M1	3
M2	4
M3	5

These are not minimal but they are small and thus suffice to prove M1, M2 and M3 tractable, because M1–3 are special cases of MPE. Note that including the removed parameters  $\beta$ ,  $\gamma$ ,  $\kappa$  and  $w$  would increase the treewidth, but it would still be constant so the tractability result is also valid for the original model.

**Corollary 3.1.** Because M1, M2 and M3 have treewidth  $\leq 5$ , M1, M2 and M3 are tractable.

### 3.2 MULTIPLE GOALS BIP

The tractability of M1–3 is in some sense an artifact of the simplified experiments for which these models were designed. Baker et al. (2009) criticize their own model by explaining their assumption of complete observability is unrealistic. One could propose to break that assumption, introduction believe variables, less observed variables or both. This generalization can impact the scalability and tractability of the model.

Another assumption in the BIP model are simplistic state and action representations. Complex states and actions require many values to be encoded, this is psychologically implausible and complexity analysis does not allow exponentially growing encodings. Breaking this assumption would introduce more state and action variables per time step, including their dependencies.

A third assumption is that, under the BIP model an observer can not assume an agent has more than one 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, unless they are encoded in the goal values. Doing so would require an exponential number of values, exponentially increasing the size of the encoding which is both unrealistic from a cognitive perspective and forbidden in complexity analysis.

Other extensions are possible as well. In (Ullman, Baker, Macindoe, Goodman, & Tenenbaum, 2009) the authors extend the original BIP model to describe goal inference in situations where the observer tries to help the agent.

In this thesis the third assumption is broken to demonstrate how complexity analysis can be used to analyse (intractable) cognitive models. To accommodate for goal inferences where multiple goals are inferred, we propose an extension called MULTIPLE GOALS BIP or MGBIP. Fig. 3.5 illustrates the dynamic Bayesian network of MGBIP. There are multiple sets of goal variables  $\mathbf{G}_1 \dots \mathbf{G}_k$ , each action  $A_t$  depends in some way on any of the variables in the sets. In the MGBIP model we call a set of goal variables a *multiple goal*.

### 3.3 MGBIP is intractable

Because it is more general, MGBIP has wider range of applicability than M1–3 but the introduced generality also comes at a cost. Whereas M1, M2 and M3 are tractable MGBIP is intractable: there are no tractable –






Figure 3.5: Graphical representation of the dynamic Bayesian network that describes MULTIPLE GOALS BIP (MGBIP).

---

polynomial time – algorithms that can implement this model.

Shimony (1994) proved finding MPE is *NP*-hard in general BNs. We show that, even while it is a special case of MPE with restricted topology, MULTIPLE GOALS BIP (MGBIP) is also *NP*-hard. To prove MGBIP is intractable, we provide a polynomial time reduction from DECISION-3SAT to DECISION-MGBIP and we argue that because DECISION-MGBIP is *NP*-hard, MGBIP is intractable. 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 in  $U$ . Each clause is a disjunction of at most three variables.

**Output:** Does there exist 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 3.5)  $\mathcal{B} = (\mathbf{G}_1, \dots, \mathbf{G}_k, \mathcal{D})$  where  $k > 0$ , 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  $\Pr(G = \mathbf{g}) \geq q$ ?

To rewrite a D-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 3.1.** *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  $\{X_1, \dots, X_k\}$ , where each of the variables can be negated. The negations are encoded in the conditional probability of the clause. Let  $\neg_p$  be true if and only if the  $p^{th}$  position of the clause is negated. We define the conditional probability of the clause variable as:

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

Clause variable probability distributions for clauses with less variables can be defined analogous.

**Lemma 3.1.** D-MGBIP is NP-hard.

In the proof we degrade dependencies in the BIPBN. To define a degraded dependency let  $C$  depend on  $A$  and  $B$ . Suppose we have to provide the conditional probabilities for the BN and each variable can assume either *true* or *false*. Then we need to provide the following conditional probabilities:

$$\begin{aligned} \Pr(C = \text{true} \mid A = \text{true}, B = \text{true}) &= \alpha \\ \Pr(C = \text{true} \mid A = \text{true}, B = \text{false}) &= \beta \\ \Pr(C = \text{true} \mid A = \text{false}, B = \text{true}) &= \gamma \\ \Pr(C = \text{true} \mid A = \text{false}, B = \text{false}) &= \delta \end{aligned}$$

If we set  $\alpha = \beta$  and  $\gamma = \delta$ , then it does not matter what evidence we have for  $B$ . The conditional probability of  $\Pr(C \mid B)$  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  $\varphi$  to an instance of MGBIP  $\mathcal{B}$ , we create a multiple goal  $\mathbf{G}_i$  containing one goal variable  $G_i$  for each variable in  $\varphi$ . For each clause in  $\varphi$  an action with the corresponding clause probability distribution is created in  $\mathcal{B}$  and for each conjunction in  $\varphi$  we create a conjunction node at state  $S_{t+1}$ , its conditional probability  $\Pr(S_{t+1} | S_t, A_t) = 1$  if  $S_t = true$  and  $A_t = true$  and 0 otherwise. Furthermore we set  $S_0 = true$  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  $\mathbf{G}_1, \dots, \mathbf{G}_k$  for which  $\Pr(\mathbf{g}) \geq q$ . All dependencies between a goal node and an 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 3.6 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, satisfying the criteria for a polynomial time reduction:

1. If  $\varphi$  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 3.1 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  $\mathbf{G}_1, \dots, \mathbf{G}_k$  is 0 if it does not satisfy all clauses, or 1 if it does.
2. If  $\mathcal{B}$  is a yes-instance, then  $\varphi$  is a yes-instance: Given the conditional probability  $\Pr(S_{t+1} | S_t, A_t)$ ,  $\mathbf{G}_1, \dots, \mathbf{G}_k$  need to be consistent with each clause variable in the BN. If  $\mathcal{B}$  is a yes-instance then  $\Pr(\mathbf{g}) = 1$  and the joint value assignment  $\mathbf{g}$  for  $\mathbf{G}_1, \dots, \mathbf{G}_k$  is consistent with each clause variable. Per definition of the clause variable's conditional probability distribution value assignment  $\mathbf{g}$  satisfies each clause in  $\varphi$ .
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 3.2.** If D-MGBIP is NP-hard then MGBIP is intractable.

*Proof.* Assume there exists a polynomial time algorithm that solves MGBIP (viz. it returns the most probable explanation  $\mathbf{g}$  for  $\mathbf{G}_1, \dots, \mathbf{G}_k$ ). Then together with the observations  $\mathbf{s}$  for  $\mathbf{S}$  and  $\mathbf{a}$  for  $\mathbf{A}$  we can compute  $\Pr(\mathbf{g} \mid \mathbf{s}, \mathbf{a})$  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 NP-hard, thus we have an inconsistency and we reject that mGBIP is solvable in polynomial time. □

**Corollary 3.2.** MGBIP is intractable, because D-MGBIP is NP-hard.




Figure 3.6: 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<sub>0</sub> and clause<sub>1</sub> are:  $\Pr(C_0 = true \mid G_1 = true \vee G_2 = false \vee G_3 = false) = 1$  and 0 otherwise,  $\Pr(C_1 = true \mid G_1 = true \vee G_2 = true \vee G_3 = false) = 1$  and 0 otherwise.

---

Proving MGBIP is intractable contradicts the fact that 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 chapter describes how we set out to identify such possible constraints.



# 4

## 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. One starts by identifying a set of model *parameters*  $\kappa$  in the model  $M$  under study (for us, MGBIP), that are possible sources of intractability. 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  $\kappa$  yet slowly (polynomial) in the size of the input, i.e. if  $\kappa$ - $M$  is fp-tractable.

The MGBIP model has several parameters, each of them a candidate source of intractability. In this paper we consider five such parameters that – on intuitive grounds – may be considered candidate sources of intractability in the MGBIP model (see Table 4.2 for an overview and Fig. 4.1 for an illustration).

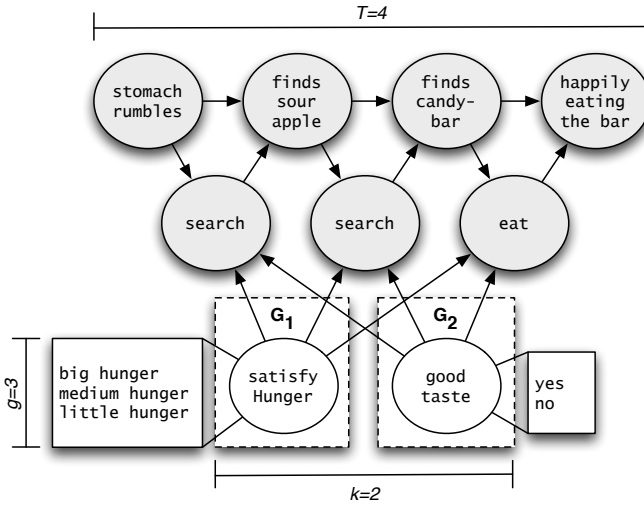


Figure 4.1: Illustration of the Bayesian network and different parameters of the MGBIP model applied to the “mother observes son”-example.

Table 4.1: Example probability distribution over the combinations of goal values. In this example  $p = 0.6$  and  $1 - p = 0.4$ .

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



Table 4.2: 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
$s$	maximum # state values	4
$a$	maximum # action values	2
$g$	maximum # goal values	3
$1 - p$	distance from certainty	0.4

## 4.1 MGBIP fp-intractability

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, the parameters  $s$ ,  $a$  and  $g$  are the maximum number of values per – respectively – each state, action and goal variable. As the number of possible values that a variable can take increases the necessary number of calculations, also  $s$ ,  $a$  and  $g$  is a candidate source of intractability.

Based on these parameters we prove MGBIP is fp-intractable for every subset of parameters  $\kappa \subseteq \{T, 1/T, g\}$ , contradicting the intuition about their role in the tractability of the model. Note that based on Lemma 2.3 it is required to prove MGBIP intractable both for  $T$  and  $1/T$ .

**Proposition 4.1.** MGBIP is not fixed-parameter tractable for  $\{s, a, g\}$ .

*Proof.* The NP-hardness proof of MGBIP (Section 3.3) 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.  $\square$

**Proposition 4.2.** 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 small.  $\square$

**Proposition 4.3.** MGBIP is fp-intractable for  $\{1/T\}$ .

*Proof.* If the reduction from D-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  $\Pr(S'_t \mid S_{t-1} = \text{true}, A_{t-1} = \text{true}) = 1$  or 0 otherwise and the conditional probability  $\Pr(A'_t \mid G_i = g_i, \dots, G_j = g_j, S_{t-1} = \text{true}) = 1$  or 0 otherwise, where  $g_i \dots g_j$  can be any value (i.e.  $A'_t$  is independent of all goals). This means we can reduce any D-3SAT instance to D-MGBIP even when  $1/T$  is small.  $\square$

**Proposition 4.4.** 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 Proposition 4.2 and Proposition 4.3, thus we can conclude that such an algorithm does not exist.  $\square$

Because the proofs of propositions 4.1-4.4 do not assume more than two values for each variable and because lemma 2.2 we observe:

**Result 4.1.** MGBIP is fp-intractable for every subset of parameters  $\kappa \subseteq \{T, 1/T, g\}$ .

Result 4.1 shows – contrary to the intuitions – 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.

## 4.2 MGBIP fp-tractability

Now consider parameter  $k$ , 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.

**Proposition 4.5.** MGBIP is fp-tractable for  $\{k\}$ .

*Proof.* We know that MPE is fixed-parameter tractable for treewidth (J. Kwisthout, 2010) and MGBIP is a special case of MPE (Section 3.3). 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.  $\square$

**Result 4.2.** MGBIP is fp-tractable for parameter  $\{k\}$ .

Result 4.2 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 observers can only (quickly) infer multiple goals if agents they observe pursue only a small number of goals. This seems

realistic because either agents do not (typically) pursue a large number of goals in parallel at the same time (possibly also to keep their own planning tractable).

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. See e.g. Table 4.1, where  $1 - p = 0.4$  is relatively small and the most likely explanation satisfy hunger=little hunger, desire sweet=yes has little competition. If the value is large, it means that the most likely explanation has many competitor explanations of non-negligible probability. 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.

**Proposition 4.6.** MGBIP is fp-tractable  $\{1 - p\}$ .

*Proof.* It is known that MPE is fixed-parameter tractable for probability of the most probable explanation (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 of the most probable explanation.  $\square$

**Result 4.3.** MGBIP is fp-tractable for parameter  $\{1 - p\}$ .

Result 4.3 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.



# 5

## 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 systematic behavior. Recall, for example, the scenario from the Introduction where the son searches the kitchen for a candy bar. Under different circumstances, the

mother may understand that her son has even more goals at a single point in time, for example: to still his hunger, to satisfy his craving for sweet, to see how many bars are left, to pretend that he did not hear his mom's request to clean up his room, to bring back a candy bar for his mom, 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 (see 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, action or state 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 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, Gergely, Biró, Koós, and Brockbank (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: meaning subjects will perform either slow or make bad inferences. 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 any Bayesian model, whether of action understanding or otherwise. The approach reveals that – contrary to popular belief in cognitive psychology – 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.



# 6

## Bibliography

### 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.
- Baker, C. L., Tenenbaum, J., & Saxe, R. (2007, Jan). Goal inference as inverse planning. *Proceedings of the 29th meeting of the Cognitive Science Society*.
- Baldwin, D. A., & Baird, J. A. (2001). Discerning intentions in dynamic human action. *Trends in Cognitive Sciences*, 5(4), 171–178.
- 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.
- 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., Erhagen, W., & Bekkering, H. (2006, Jan). Goals and means in action observation: A computational approach. *Neural Networks*.
- Downey, R., & Fellows, M. (1998). *Parameterized complexity* (Heidelberg, Ed.). Springer-Verlag.
- Garey, M. R., & Johnson, D. S. (1979). *Computers and intractability: A guide to the theory of np-completeness (series of books in the mathematical sciences)*. W. H. Freeman.
- Gergely, G., Nádasdy, Z., Csibra, G., & Biró, S. (1995). Taking the intentional stance at 12 months of age. *Cognition*, 56, 165–193.
- Ghahramani, Z. (1998). Learning dynamic Bayesian networks. In *Adaptive processing of sequences and data structures* (pp. 168–197). Springer-Verlag.
- 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.
- Hassin, R., Aarts, H., & Ferguson, M. (2005, Jan). Automatic goal inferences. *Journal of Experimental Social Psychology*.
- Jensen, F. V., & Nielsen, T. D. (2007). *Bayesian networks and decision graphs* (Second ed.). New York: Springer Verlag.
- Király, I., Jovanovic, B., & Prinz, W. (2003, Jan). The early origins of goal attribution in infancy. *Consciousness and Cognition*.
- Kwisthout, J. (2010). *Most probable explanations in bayesian networks: complexity and algorithms*.
- Kwisthout, J. H. P. (2009). *The computational complexity of probabilistic networks*. Unpublished doctoral dissertation, Faculty of Science, Utrecht University, The Netherlands.
- Marr, D. (1982). *Vision : a computational investigation into the human representation and processing of visual information / david marr* [Book]. W.H. Freeman, San Francisco ; New York :.

- 
- 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.
- 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.
- 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.
- Sipser, M. (1992). The history and status of the p versus np question. In *Stoc '92: Proceedings of the twenty-fourth annual acm symposium on theory of computing* (pp. 603–618). New York, NY, USA: ACM.
- 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).
- 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., 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.