

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

The computational costs of recipient design and intention recognition in communication

Permalink

<https://escholarship.org/uc/item/56n6s4g6>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 33(33)

ISSN

1069-7977

Authors

Blokpoel, Mark
Kwisthout, Johan
Wareham, Todd
et al.

Publication Date

2011

Peer reviewed

The computational costs of recipient design and intention recognition in communication

¹Mark Blokpoel (m.blokpoel@donders.ru.nl), ²Johan Kwisthout

³Todd Wareham, ¹Pim Haselager, ¹Ivan Toni, ¹Iris van Rooij

¹Radboud University Nijmegen, Donders Institute for Brain Cognition and Behaviour, The Netherlands

²Radboud University Nijmegen, Institute for Computing and Information Sciences, The Netherlands

³Department of Computer Science, Memorial University of Newfoundland, Canada

Abstract

Understanding the communicative intentions of others based on their behavior can be seen as an ‘inference to the best explanation’, a.k.a. abduction. As abduction is often an intractable task, it has been suggested that communicators alleviate the work of an addressee by performing recipient design, adapting their behavior to the presumed beliefs and knowledge of the addressee. In this paper we show that communicators performing recipient design inherit the computational load of their addressees. Thus, recipient design in itself cannot explain the speed of everyday human intentional communication.

Keywords: Bayesian network; communication; recipient design; simulation; abduction; inference; NP-hard; tractable.

Introduction

Humans have the ability to understand the intentions underlying communicative actions of others. This is a remarkable ability given that intention recognition involves reasoning from effects (observed actions) to their likely causes (hypothesized intentions) and is therefore best seen as a form of ‘inference to the best explanation’, a.k.a. *abduction* (Levinson, 2006; Sperber & Wilson, 1995; Baker, Saxe, & Tenenbaum, 2009; Charniak & Goldman, 1991; Peirce, 1931–1966; Lipton, 2004). Computational models of abduction are notorious for their computational intractability, meaning that the inferences postulated by these models require exponential amounts of time (Abdelbar & Hedetniemi, 1998; Thagard & Verbeurgt, 1998; Kwisthout, 2010; Blokpoel, Kwisthout, van der Weide, & van Rooij, 2010; Bylander, Allemang, & Tanner, 1991; Nordh & Zanuttini, 2005). Evidently, intractable models cannot explain the speed of intention recognition as we observe in everyday communication.

It has been suggested that the computational demands of intention recognition in human communication could be alleviated through recipient design, in which communicative actions are constructed according to what addressees are supposed to know and believe (see Box 1; Sperber and Wilson (1995); Grice (1989); Clark (1996)). This idea is generally consistent with theoretical work showing that intention recognition can be tractable given specific constraints (Blokpoel et al., 2010), and with empirical work qualifying the conditions under which recipient design is used (Clark, 1996; Keysar, Barr, & Horton, 1998; Newman-Norlund et al., 2009). However, the idea also raises a so far neglected question: If recipient design is assumed to make intention recognition tractable for addressees, does it not simply move the computational load from the addressee to the agent generating the communicative action?

In this paper we present a formal model of communication and prove that even under highly restricted conditions recipient design is intractable. This proof of intractability of recipient design establishes that even though recipient design can make intention recognition tractable, the computational demands of recipient design are such that the speed of everyday communication is not yet explained. This finding indicates that communicators must exploit constraints to make recipient design tractable, and in the second part of the paper we illustrate a methodology suitable for identifying such constraints.

Computational-level Models

In this section we present a formal computational-level model of communication based on the Bayesian Inverse Planning (BIP) model of action understanding by Baker et al. (2009) and on the statistical learning model by Shafto and Goodman (2008). Empirical evidence presented by these authors shows that these models seem to capture fundamental principles underlying intention recognition and recipient design respectively. Our model of communication combines these two models and will as a result inherit some of their simplifying assumptions. Consequently, our analyses will yield at worst a lower bound of the computational demands posited by more general models of communication.

The communication model we present assumes that a communicator generates communicative behavior by choosing actions to achieve certain goals. These goals can be divided in two types: instrumental (e.g. ‘make the mosquito go away’) and communicative (e.g. ‘signal the taxi driver to come here’). Because some actions can lead to the achievement of more than one (type of) goal (e.g. ‘waving one’s hand’ can make a mosquito go away, but also signal a taxi driver), recognizing communicative intentions involves abduction. Furthermore communicators also perform recipient design, choosing their actions on the basis of world states, instrumental goals and communicative goals, but also on the basis of a prediction of the likely inferences their audience could make given the action sequence they intend to produce.


These characterizations of RECIPIENT DESIGN and INTENTION RECOGNITION can be summarized by input-output mappings. A communicator generates a *sequence of actions* that will (a) most likely lead to the achievement of the instrumental goals and (b) will lead his/her audience to attribute the correct communicative goals to the communicator’s behav-

Box 1: An illustration of recipient design.

To illustrate recipient design, consider the following example. Bob and Mary are chatting, while suddenly Bob's favorite composition by Bach sounds faintly in the background. Bob wants to communicate three things to Mary:

- i. He wishes her to be quiet;
- ii. He wants her to listen to the music;
- iii. He wants to signal he is listening to the music.

Now suppose that Bob knows that Mary knows they both enjoy Bach very much. To communicate (i), (ii) and (iii) to her he might simply just put his finger in front of his mouth (a).



Placed in a different situation talking to Ann, Bob might communicate differently. He knows Ann likes to keep talking and that she is not interested in music. He put his finger in front of his mouth to signal Ann to be quiet (i), but he also closes his eyes to tell her he is listening (iii) (b). To emphasize he is listening (iii) even more, Bob then tilts his head slightly and puts his finger up in the air, signaling Ann to pay attention and listen to the music (ii) (c). [Illustrations by Bas Maes.]

ior. This inference is based on the probabilistic dependencies between actions and world states (including how these dependencies change over time) and zero or more instrumental goals and one or more communicative goals. The addressee infers a *combination of communicative goals* that best explains the observed communicative behavior given what he/she knows about the probabilistic dependencies between actions, goals and world states (including how these dependencies change over time).

We define the following variables that we use to formalize the input-output mappings for RECIPIENT DESIGN and INTENTION RECOGNITION.¹

- $\mathbf{S} = \{S_1, \dots, S_T\}$, a sequence of T state variables that can encode values of state sequences \mathbf{s} ;
- $\mathbf{A} = \{A_1, \dots, A_{T-1}\}$, a sequence of $T - 1$ action variables that can encode values of action sequences \mathbf{a} ;
- $\mathbf{G}_I = \{G_{I1}, \dots, G_{Ij}\}$, a set of instrumental goal variables that can encode the values of the communicator's instrumental goals \mathbf{g}_I ; and
- $\mathbf{G}_C = \{G_{C1}, \dots, G_{Ck}\}$, a set of communicative goal variables that can encode the values of the communicator's communicative goals \mathbf{g}_C .

RECIPIENT DESIGN

Input: A Bayesian network $\mathcal{B} = (N, \Gamma)$, a value assignment \mathbf{g}_I for \mathbf{G}_I and a value assignment \mathbf{g}_C for \mathbf{G}_C encoding the communicator's goals.

Where, $\mathbf{S}, \mathbf{A}, \mathbf{G}_I, \mathbf{G}_C \in N$; the probabilistic dependencies in N are illustrated in Figure 1; and Γ is an arbitrary conditional probability distribution over N .

Output: A value assignment \mathbf{a} to \mathbf{A} , such that $\mathbf{a} = \operatorname{argmax}_{\mathbf{a}} \Pr(\mathbf{A} = \mathbf{a} \mid \mathbf{G}_I = \mathbf{g}_I)$ and $\text{INTENTION RECOGNITION}(\mathcal{B}, \mathbf{a}, \mathbf{s}) = \mathbf{g}_C$, or \emptyset if no sequence of actions \mathbf{a} is possible. Here $\mathbf{s} = \operatorname{argmax}_{\mathbf{s}} \Pr(\mathbf{S} = \mathbf{s} \mid \mathbf{A} = \mathbf{a})$, i.e. the most likely states \mathbf{s} to follow from the actions.

INTENTION RECOGNITION

Input: A Bayesian network $\mathcal{B} = (N, \Gamma)$, similar as in the Recipient Design network, a value assignment \mathbf{a} for \mathbf{A} and a value assignment \mathbf{s} for \mathbf{S} encoding the observed actions and states.

Output: The most probable value assignment \mathbf{g}_C to the communicative goals \mathbf{G}_C , i.e. $\operatorname{argmax}_{\mathbf{g}_C} \Pr(\mathbf{G}_C = \mathbf{g}_C \mid \mathbf{S} = \mathbf{s}, \mathbf{A} = \mathbf{a})$, or \emptyset if $\Pr(\mathbf{G}_C = \mathbf{g}_C \mid \mathbf{S} = \mathbf{s}, \mathbf{A} = \mathbf{a}) = 0$ for all possible values for \mathbf{G}_C .

Recipient Design is Intractable

To investigate the computational (in)tractability of RECIPIENT DESIGN we adopted complexity-theoretic proof techniques (see e.g. Garey and Johnson (1979)). Using these

¹In the Bayesian formalism capital letters denote variables, whereas small letters denote values; bold letters denote sets of variables or values, whereas non-bold letters denote singletons.




Figure 1: The Bayesian network showing the dependencies between the variables in the RECIPIENT DESIGN and INTENTION RECOGNITION models. Arrows denote dependencies, e.g. if Bob has his eyes open ($S_t = \text{'Bob eyes open'}$), then closes his eyes ($A_t = \text{'close eyes'}$), then S_{t+1} has a high probability of Bob having his eyes closed ($\Pr(S_{t+1} = \text{'Bob eyes closed'} | S_t = \text{'Bob eyes open'}, A_t = \text{'close eyes'}) = 0.9$).

techniques, we proved the following (see online supplementary materials for the full proofs²):

Result 1. RECIPIENT DESIGN is NP-hard.

This result implies that there does not exist any algorithm that can compute the recipient design input-output function in polynomial time for all its inputs (i.e., a time upper bounded by some function n^c where n is a measure of input size and c is some constant).³ In other words, all algorithms solving RECIPIENT DESIGN will run in exponential time or worse for a non-empty set of inputs (i.e., a time at best upper bounded by some function c^n , where n is again a measure of input size and c a constant). As exponential time algorithms run unrealistically long for all but very small inputs they are generally considered computationally intractable (Garey & Johnson, 1979). To illustrate this point, consider an exponential-time algorithm running in a time proportional to 2^n . Such an algorithm would need to make on the order of 1,000,000,000 computational steps for an input of size $n = 40$, which is more milliseconds than there are in a millennium.

Our NP-hardness result is quite sobering, given that the recipient design model already incorporates several simplifying assumptions. For instance, the model assumes communicators have perfect (probabilistic) knowledge of the world and the audience; states and goals are probabilistically independent; and there is no higher-order reasoning by communicator and audience about each other's beliefs (Verbrugge, 2009; Shafto & Goodman, 2008). This means that Result 1 probably underestimates the computational complexity of recipient design under less restricted conditions—i.e., other more general models of recipient design may well be computationally

²<http://www.dcc.ru.nl/~irisvr/suppl2011.pdf>

³Our interpretation assumes that the $P \neq NP$ conjecture is true. This mathematical conjecture is unproven to date, but widely believed by mathematicians on both theoretical and empirical grounds (Fortnow, 2009; Garey & Johnson, 1979).

even more demanding than the simplified RECIPIENT DESIGN function.

Though Result 1 serves to illustrate the non-trivial nature of explaining the speed of communication, we certainly do not wish to suggest that it is in principle impossible to explain the speed of communication in everyday life. Result 1 merely establishes that a computational explanation of the speed of communication will require that one incorporates one or more explicit hypotheses about situational constraints that make the (otherwise intractable) recipient design task performed by a communicator tractable. In the next section we present and illustrate a methodology that communication researchers may adopt to model and test the validity of such constraints.

A Method for Identifying Tractability Conditions

In order to find constraints on the input domain of RECIPIENT DESIGN that render the (restricted) model tractable, we adopt methods derived from parameterized complexity theory (Downey & Fellows, 1999; van Rooij & Wareham, 2008). Parameterized complexity theory is an extension of classical complexity theory motivated by the observation that it is sometimes possible that an NP-hard function $M : I \rightarrow O$ can be computed by algorithms whose running time is polynomial in the overall input size n and non-polynomial only in some aspects of the input called input parameters. In other words, the main part of the input contributes to the overall complexity in a “good” way, whereas only the input parameters contribute to the overall complexity in a “bad” way. In such cases, the function M is fixed-parameter tractable for that respective set of parameters. The following definition states this idea more formally.

Definition 1. *Fixed-parameter (fp-) tractability.* Let $M : I \rightarrow O$ be an input-output function with input parameters k_1, k_2, \dots, k_m . Then M is said to be fixed-parameter tractable for parameter-set $K = \{k_1, k_2, \dots, k_m\}$ if there exists at least one algorithm that computes O for any input of size n in time $f(k_1, k_2, \dots, k_m)n^c$, where $f(\cdot)$ is an arbitrary computable function and c is a constant. If no such algorithm exists then M is said to be fixed-parameter intractable for parameter-set K .

Note that if an intractable function M is fp-tractable for parameter-set K , then M can be efficiently computed even for large inputs, provided only that all the parameters in K are small. This means that if M is postulated as an explanation of the functional form of the input-output mapping computed by a given process, then the speed of that process in certain situations can be explained by postulating that the parameters in K are small exactly in those situations (see also van van Rooij and Wareham (2008)). This strategy for rendering (otherwise intractable) theories tractable has been successfully applied in various domains (van Rooij, Evans, Müller, Gedge, & Wareham, 2008; Müller, van Rooij, & Wareham, 2009; Wareham, Evans, & van Rooij, 2010; van Rooij, 2008; van Rooij, Stege, & Kadlec, 2005), including the Bayesian Inverse Planning

model (Blokpoel et al., 2010). In the next section we report on our investigation of parameters that do and do not render RECIPIENT DESIGN tractable.

What Makes Recipient Design Tractable?

The RECIPIENT DESIGN model has several parameters that we will consider for our fixed-parameter (fp-)tractability analyses. Table 1 gives an overview of these parameters and their example values in the illustration in Figure 2. Proofs of all these results can be found in the Supplementary Materials published online.⁴

Parameter	Description	Value
$ \mathbf{G}_C $	The number of communicative goals	3
$ \mathbf{G}_I $	The number of instrumental goals	0
$ \mathbf{A} $	The number of observed or planned actions	2

Table 1: Overview of the parameters, the given value is based on the Bob-Ann example in Box 1.




Figure 2: An example network with all values filled in. Here Bob would have to find actions given his communicative goals and Ann would have to infer Bob’s communicative goals given his actions and states.

We start by considering conditions that render intention recognition tractable. The following results are relevant for our purposes.

Result 2. INTENTION RECOGNITION is NP-hard.

Result 3. INTENTION RECOGNITION is fp-intractable for parameter sets $\{|\mathbf{A}|, |\mathbf{G}_C|\}$ and $\{|\mathbf{A}|, |\mathbf{G}_I|\}$.

Result 4. INTENTION RECOGNITION fp-tractable for the parameter set $\{|\mathbf{G}_I|, |\mathbf{G}_C|\}$.

For an overview of further fp-(in)tractability results implicated by Results 3 and 4, see Table 2.

Result 2 establishes that without any constraints on the input domain, INTENTION RECOGNITION is intractable—just

as RECIPIENT DESIGN (Result 1)—in that its computation requires superpolynomial time. Result 4 shows that INTENTION RECOGNITION can be computed efficiently even for large inputs provided only that two parameters $|\mathbf{G}_I|$ and $|\mathbf{G}_C|$ are both relatively small. As both these parameters seem to be under the control of the communicator, Result 4 presents the first formal explication of the hypothesis that a communicator may make the task of the audience to infer his/her intentions easier and even tractable.

Note furthermore that relative to the parameters that we consider (i.e., $\{|\mathbf{G}_I|, |\mathbf{G}_C|, |\mathbf{A}|\}$), Result 3 and 4 combined show that the parameter set $\{|\mathbf{G}_I|, |\mathbf{G}_C|\}$ is not only sufficient but also necessary for fp-tractability. That is, INTENTION RECOGNITION is fp-intractable for all proper subsets of the $\{|\mathbf{G}_I|, |\mathbf{G}_C|, |\mathbf{A}|\}$ and for other subsets that do not include $\{|\mathbf{G}_I|, |\mathbf{G}_C|\}$.

Having identified constraints that a communicator may utilize to render INTENTION RECOGNITION tractable, a natural question to ask is whether recipient design is tractable under these same constraints. The following result shows this is not the case.

Result 5. RECIPIENT DESIGN is fp-intractable for the parameter sets $\{|\mathbf{G}_I|, |\mathbf{G}_C|\}$, $\{|\mathbf{G}_I|, |\mathbf{A}|\}$ and $\{|\mathbf{G}_C|, |\mathbf{A}|\}$.

Result 5 shows that RECIPIENT DESIGN is strictly more difficult than INTENTION RECOGNITION, as the former is not tractable under conditions that make the latter tractable. It also means that the computational intractability of recipient design cannot be attributed solely to the complexity of simulating the audience’s intention recognition processes as a subroutine.

We conclude with the following result:

Result 6. RECIPIENT DESIGN is fp-tractable for the parameter set $\{|\mathbf{A}|, |\mathbf{G}_I|, |\mathbf{G}_C|\}$.

Result 6 shows that RECIPIENT DESIGN can be computed efficiently provided that all three parameters $|\mathbf{G}_I|$, $|\mathbf{G}_C|$ and $|\mathbf{A}|$ are relatively small. Note, however, that restricting all three parameters at the same time effectively ensures the whole input network is small, and hence the parameters cannot figure in an explanation of how communication can be tractable for large input networks. As shown in Table 2, no proper subset of $\{|\mathbf{A}|, |\mathbf{G}_I|, |\mathbf{G}_C|\}$ suffices to make RECIPIENT DESIGN fp-tractable. Although other parameters than the ones considered here may figure in an explanation of the speed of communication, our findings highlight the nontrivial problem of finding such an explanation.

Discussion

As with many other core human abilities, intention recognition appears a fairly straightforward phenomenon, at least until we interact with other humans having communication deficits, or until we try to build artificial cognitive agents that can effectively implement flexible intention recognition in a communicative setting. The astronomical computational powers required for abductive processes such as intention

⁴<http://www.dcc.ru.nl/~irisvr/suppl2011.pdf>

INTENTION RECOGNITION	—	A
—	NP-hard	fp-intract.
G _C	fp-intract.	fp-intract.
G _I	fp-intract	fp-intract
G _C , G _I	fp-tractable	fp-tractable
RECIPIENT DESIGN	—	A
—	NP-hard	fp-intract.
G _C	fp-intract.	fp-intract.
G _I	fp-intract.	fp-intract.
G _C , G _I	fp-intract.	fp-tractable

Table 2: Complexity results for INTENTION RECOGNITION (above) and RECIPIENT DESIGN (below).

recognition are in contrast with the speed of everyday communication. To explain this contrast, it has been suggested that intention recognition may be made easier if communicators use recipient design (Sperber & Wilson, 1995; Grice, 1989). The aim of our research was to assess to what extent this idea merely shifts the computational complexity of communication from the audience to the communicator. Specifically, we questioned whether computational models of recipient design inherit the computational load they aim to take away from the audience.

To address this question we formalized the tasks of communicator and audience as computational-level models. We ensured that our modeling choices did not lead to an artificial overestimation of the computational complexity of communication by incorporating several simplifying assumptions. There are two main findings. Both the audience model (i.e., INTENTION RECOGNITION) and the communicator model (i.e., RECIPIENT DESIGN) are intractable (NP-hard). This means that, notwithstanding our simplifying assumptions, the computations postulated by our models require an unrealistic amount of time for their completion.

The intractability result for INTENTION RECOGNITION reiterates what has long been assumed. Namely, given that intention recognition is a form of abduction, the speed at which we can use this ability in our everyday life is comparably difficult to explain. Replicating this result in such a simplified model underscores the non-triviality of explaining the speed of intention recognition. The main novelty of this study lies in defining the computational demands of recipient design, an undeservedly overlooked issue given the centrality of this ability to several accounts of communication (Sperber & Wilson, 1995; Clark, 1996). The intractability result for RECIPIENT DESIGN shows that even if communicators can make intention recognition easier by performing recipient design, the model by itself cannot explain the speed of every day communication. These results set the stage for both theoretical and empirical follow-up research.

From a theoretical perspective, the intractability results raise the question how the speed of everyday communication can be reconciled with the apparent complexity of the tasks performed by communicator and audience. This question can be addressed by identifying the situational constraints that

render the tasks of communicator and audience tractable. We have presented a methodology for implementing this strategy and illustrated its use for our models. We found that if the communicator has only a few communicative and instrumental goals, INTENTION RECOGNITION is tractable. These special circumstances are, however, not yet sufficient to also make RECIPIENT DESIGN tractable. The additional circumstance where a communicator is able to construct short action sequences to convey his/her message does make RECIPIENT DESIGN tractable. In other words, under the simplifying assumptions of our models, people might exploit these special conditions to achieve speedy communication.

These conditions may not suffice to explain the speed of communication in general, since some of our simplifying assumptions most probably will be violated in real world situations. Yet this underscores that richer models of recipient design—with less simplifying assumptions, e.g., including higher-order reasoning—will presumably be even more computationally demanding. Therefore richer models would also require an analysis of their computational demands.

The utility of the current approach can also be assessed empirically by creating experimental set-ups which do meet the simplifying conditions (Galantucci, 2005; de Ruiter et al., 2010; Scott-Phillips & Kirby, 2010). In such experimental set-ups it can then be tested if the constraints that we identified as necessary and sufficient for tractability of communication are confirmed by the success or failure of communication as observed in the lab.

It might be relevant to emphasize that the present results converge with several intuitions of classic pragmatic theories such as the Gricean Maxims (Grice, 1989). For example, the Maxim of Quantity states that people should not make their contribution more informative than is required. In the current models, “informativeness” could be operationalized as the number of communicative goals a communicator tries to convey. The Maxim of Quantity can then be interpreted as not having too many communicative goals, which is equivalent to one of the constraints necessary for tractability of the communication models. Grice’s Maxim of Relation states people should be relevant. Relevance in our models can be indexed by the number of instrumental goals that influence one’s communicative behavior. Having few instrumental goals increases the communicative relevance of the communicator’s behavior, making it easier for an audience to perform intention recognition. This principle is similar to the necessary constraint of pursuing few instrumental goals highlighted by the current tractability analysis.

The strong convergence between Grice’s Maxims and the current results suggests that the communicator and audience models capture at least some fundamental aspects of communication and recipient design. It also suggests that the current approach could provide a formal account of the cognitive mechanisms described by those maxims, enabling more systematic empirical analyses.

To conclude, we showed that by performing recipient de-

sign, a communicator may reduce the computational load of her addressees, but this then leaves the communicator facing an intractable task. The fact that this result is based on highly simplified models greatly underscores the non-triviality of explaining the speed of everyday communication, as more general models will also suffer from intractability. This result highlights an explanatory gap in communication science, and we illustrated a methodology to deal with this gap.

Acknowledgments

JK was supported by the OCTOPUS project under the responsibility of the Embedded Systems Institute. TW was supported by NSERC Personal Discovery Grant 228104. IT was supported by grant #453-08-002 from NWO.

References

- Abdelbar, A., & Hedetniemi, S. (1998, Jun). Approximating MAPs for belief networks is NP-hard and other theorems. *Artificial Intelligence*, 102(1), 21–38.
- Baker, C., Saxe, R., & Tenenbaum, J. B. (2009, Dec). Action understanding as inverse planning. *Cognition*, 113(3), 329–349.
- Blokpoel, M., Kwisthout, J., van der Weide, T. P., & van Rooij, I. (2010). How action understanding can be rational, Bayesian and tractable. In S. O. R. Catrambone (Ed.), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society* (p. 1643-1648). Austin, TX: Cognitive Science Society.
- Bylander, T., Allemang, D., & Tanner, M. (1991, Jan). The computational complexity of abduction. *Artificial Intelligence*, 49(1–3), 25–60.
- Charniak, E., & Goldman, R. (1991, Jan). *Probabilistic abduction for plan recognition* (Tech. Rep.). Department of Computer Science Brown University and Department of Computer Science Tulane University.
- Clark, H. H. (1996). *Using language*. Cambridge: Cambridge University Press.
- de Ruiter, J., Noordzij, M., Newman-Norlund, S., Newman-Norlund, R., Hagoort, P., Levinson, S., et al. (2010). Exploring the cognitive infrastructure of communication. *Interaction Studies*, 11, 51–77.
- Downey, R. G., & Fellows, M. R. (1999). *Parameterized Complexity*. Springer. Hardcover.
- Fortnow, L. (2009). The status of the P versus NP problem. *Communications of the ACM*, 52(9), 78–86.
- Galantucci, B. (2005). An experimental study of the emergence of human communication systems. *Cognitive Science*, 29, 737–767.
- Garey, M. R., & Johnson, D. S. (1979). *Computers and intractability: A guide to the theory of NP-completeness*. W. H. Freeman.
- Grice, P. (1989). *Studies in the way of words*. Harvard University Press.
- Keysar, B., Barr, D. J., & Horton, W. S. (1998). The ego-centric basis of language use: Insights from a processing approach. *Current Directions in Psych. Science*, 7, 46–50.
- Kwisthout, J. (2010). *Most probable explanations in Bayesian networks: complexity and tractability* (Tech. Rep.). Faculty of Science, Radboud University Nijmegen.
- Levinson, S. (2006). On the human 'interaction engine'. *Roots of Human Society*.
- Lipton, P. (2004). *Inference to the best explanation* (2nd ed.). Routledge.
- Müller, M., van Rooij, I., & Wareham, T. (2009). Similarity as tractable transformation. In N. A. Taatgen & H. van Rijn (Eds.), *Proceedings of the 31st Annual Conference of the Cognitive Science Society* (pp. 50–55). Austin, TX: Cognitive Science Society.
- Newman-Norlund, S. E., Noordzij, M. L., Newman-Norlund, R. D., Volman, I. A., de Ruiter, J. P., Hagoort, P., et al. (2009, Apr). Recipient design in tacit communication. *Cognition*, 111(1), 46–54.
- Nordh, G., & Zanuttini, B. (2005). Propositional abduction is almost always hard. In *Proceedings of the 19th International Joint Conference on Artificial Intelligence (IJCAI-2005)* (pp. 534–539). Edinburgh, Scotland, UK.
- Peirce, C. S. (1931–1966). *The collected papers of Charles S. Peirce* (Vol. 8; P. W. C. Hartshorne & A. W. Burks, Eds.). Cambridge: Harvard University Press.
- Scott-Phillips, T. C., & Kirby, S. (2010). Language evolution in the laboratory. *Trends in Cognitive Sciences*, 14(9), 411–417.
- Shafto, P., & Goodman, N. (2008, Jan). Teaching games: Statistical sampling assumptions for learning in pedagogical situations. *Proceedings of the Thirtieth Annual Conference of the Cognitive Science Society*.
- Sperber, D., & Wilson, D. (1995). *Relevance: Communication and cognition*. Blackwell Publishing.
- Thagard, P., & Verbeurgt, K. (1998, Mar). Coherence as constraint satisfaction. *Cognitive Science*, 22(1), 1–24.
- 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, Jan). Identifying sources of intractability in cognitive models: An illustration using analogical structure mapping. *Proceedings of the 30th Annual Conference of the Cognitive Science Society*.
- van Rooij, I., Stege, U., & Kadlec, H. (2005). Sources of complexity in subset choice. *Journal of Mathematical Psychology*, 49(2), 160–187.
- van Rooij, I., & Wareham, T. (2008, Nov). Parameterized complexity in cognitive modeling: Foundations, applications and opportunities. *The Computer Journal*, 51(3), 385–404.
- Verbrugge, R. (2009). Logic and social cognition: The facts matter, and so do computational models. *Journal of Philosophical Logic*, 38(6), 649–680.
- Wareham, H. T., Evans, P., & van Rooij, I. (2010). What does (and doesn't) make analogical problem solving easy? *Journal of Problem Solving*, 3(2), 30–71.