

Jos J.A. van Berkum

## Anticipating communication

---

**Jos J.A. van Berkum:** Department of Dutch, Faculty of Humanities, Utrecht University, Trans 10, 3512 JK Utrecht, The Netherlands. E-mail: j.vanberkum@uu.nl

Let's imagine you just bought some old Bill Murray movies, and you're about to ask your best friend S to come over to your place to enjoy *The life aquatic* and *Lost in translation* back to back. You know S is a big Bill Murray fan, so you phone him up:

- (1) You: Hey S, it's me.
- (2) S: Oh hi, what's up?
- (3) You: Any plans for tonight?
- (4) S: Uhh, no – but I wouldn't mind having some . . .
- (5) You: Well, I just bought *The life aquatic* and *Lost in translation*.
- (6) S: Hey, those are great, I haven't seen them in ages!
- (7) You: So, how about coming over to my place tonight?
- (8) S: . . .

There you are, in the middle of a conversation, and S is about to speak. Any predictions about what might happen next?

Well, yes. Here's a reasonable, non-exhaustive list. You're in a conversation with a cooperative speaker, so his next utterance will probably be relevant to what you were talking about. You've just asked a question, so it is very likely that you'll get a response. More specifically, you've extended an invitation, which will either be accepted or turned down. Your pre-invitations in (3) and (5), which offered ways for S to gracefully decline, were accepted and encouraged, so S will probably accept. You also know S is generally keen on such movie nights and a great fan of Bill Murray, causing you to anticipate a rather enthusiastic response. Furthermore, you're pretty sure the response will be in English, and may well include uptake words like "yes", "sure", or "great" – in fact, prior experience with S may actually suggest he'll probably start with "alrighty!". You expect him to reply rapidly, without hesitations like "Well . . .". You also know what's S's voice sounds like over the phone, as well as how his pitch tends to go up when he's excited. And you've heard that he has a cold, so it'll sound a bit nasal.

Now, as predictions go, you may have it all wrong, of course. Perhaps S won't use "alrighty!" this time. Perhaps he used Vicks inhaler during (7) and suddenly sounds less nasal. Perhaps he'll tell you that his cold is too bad to go out of the house, or that he's fed up with watching movies with you anyway. He might suddenly change the subject. He might even just hang up. But will that stop you from making those predictions?

## The predictive brain

The answer is: no. Prediction is the brain's natural mode of operation, and much of it comes for free. In fact, predicting the environment in space and time is now considered one of the brain's primary functions (Bar, 2007; 2009), and the purpose of long-term memory is precisely to support such prediction (Schacter, Addis, & Buckner, 2007; Schooler, & Anderson, 1997). This perspective is relatively new: in neuroscience, the brain has long been viewed as a fundamentally bottom-up device that hierarchically builds complex percepts by piecing together aspects of sensory input (e.g., from simple circular receptive fields to line segments to faces). But those days are over (Mesulam, 2008), and models of detailed brain function are increasingly couched in terms of prediction and minimizing prediction error (e.g., Friston, 2005; Mehta, 2001; Rao & Ballard, 1999; Summerfield, Egner, Greene, Koechlin, Mangels & Hirsch, 2006).

One of the reasons for this paradigm shift is that 'classic' hierarchical bottom-up coding cannot deliver the speedy perceptions that real brains do (Mesulam, 2008). A second, deeper reason is that the raw environmental signal is just too ambiguous and complex to deal with in a bottom-up fashion (Bar, 2007; Friston, 2005) – only predictive forward models allow us to segment, select and stabilize input so that we perceive stability and coherence in our environment (Bar, 2007). So in all, our brains are not just proactive because it pays to anticipate upcoming events in a complex dynamic world. It is also because the input would otherwise simply be too difficult to deal with efficiently.

## Smashing classic pronoun theories

The theory of pronoun interpretation proposed by Andrew Kehler and Hannah Rohde (henceforth K&R) in this issue of *Theoretical Linguistics* nicely illustrates that psycholinguistics is currently involved in a similar paradigm shift, away from bottom-up theories of language comprehension, and towards models that imple-

ment a more proactive processing style. The classic bottom-up perspective of pronoun interpretation that K&R move away from is characterized by them in terms of the ‘SMASH algorithm’: collect possible referents (Search), filter out those referents that fail ‘hard’ morpho-syntactic constraints (Match), and then select a referent on the basis of ‘soft’ constraints (And Select using Heuristics). The SMASH algorithm is essentially reactive: all processes of interest are initiated at, and by, a pronoun. It also implements a serial ‘syntax-first’ principle: morpho-syntactic constraints always have a say before pragmatic, discourse-level constraints can kick in. Actually, like many theories inspired by early digital computing, SMASH has the distinct feel of, say, a FORTRAN program running in 256 K on a big 1970’s Cyber mainframe.

Many findings, including those reviewed and newly presented by K&R, have revealed the SMASH account to be inadequate. Research with implicit causality verbs, for example, has shown that when people read “David praised Linda because . . .”, the verb and the surrounding construction leads them to expect more information about Linda, not David; if a subsequent pronoun is inconsistent with that expectation, as in “David praised Linda because he . . .”, readers slow down (Featherstone & Sturt, 2010; Koornneef & Van Berkum, 2006; Pyykkönen & Järvikivi, 2010), and also display immediate processing costs in brain potentials, right at the critical pronoun (Van Berkum, Koornneef, Otten & Nieuwland, 2007). Research has also disconfirmed the SMASH idea that ‘hard’ morpho-syntactic constraints *necessarily* apply before, and overrule, ‘soft’ discourse-level constraints. For instance, when Dutch participants read “Hij vertelde het meisje dat . . .” (“He told the girl that . . .”) in a wider story, the presence of two equally salient girls in the discourse context can momentarily lure readers into parsing “dat” as a postnominal relative pronoun heading a restrictive relative clause, rather than as a complementizer heading a complement clause. Importantly, this discourse-induced bias towards a relative pronoun reading occurs even when the grammatical gender of the noun actually *blocks* the relative pronoun reading of “dat” (as with “Hij vertelde de vrouw dat . . .”, “He told the woman that . . .”, where “dat” can still be a complementizer, but where the relative pronoun should be “die”; Van Berkum, Brown & Hagoort, 1999). Thus, discourse-level expectations can control the interpretation of “that” *before* grammatical gender filters out particular options.

The model that K&R propose is a significant step forward, in that it has proactive, large-capacity parallel processing built in as a general principle. By combining aspects of two very different views on pronoun interpretation (Hobbs, 1979; Grosz, Joshi & Weinstein, 1995) into a single parsimonious Bayesian model, K&R elegantly account for a wide range of extant findings. The model generates some non-trivial predictions, which, to the extent they are tested, bear out. And

the empirical work offered by K&R to support the model nicely exemplifies the mileage one can get with ‘simple’ fragment completion methodology – thanks to having chosen the right critical language comparisons, of course. Yet, as with any model, the one at hand leaves some things to be desired.

## Reflecting on the K&R proposal

### 1 The big question

First, there’s what K&R refer to as the BIG QUESTION. After laying out the SMASH algorithm and its complex instantiation in Centering-based pronoun interpretation models, K&R ask: “If the interpretation of pronouns is really this involved, why would any speaker do the disservice to her addressee of using one?” This is a fine moment in the paper, a pleasantly surprising turn of events, and momentarily convincing. But some more reflection led me to realize this is not really a convincing argument at all. Consider music. Open up any book on the cognition of music, and you will be struck by the complexity of the many computations that quietly do their work in the service of daily enjoyment. But should this prompt us to ask why musicians do us the disservice of creating their art? And what about the unavoidably complex computations behind visual perception, or parsing?

What is relevant here is a distinction that’s been around in cognitive science since the early days of automatic versus controlled processing experiments, and which has found its way into many modern views of the mind (e.g., Haidt, 2012; Kahneman, 2012): the distinction between automatic, effortless processing and deliberate, effortful computation. Haidt (2012) explains this in terms of an ‘elephant and rider’ metaphor, Kahneman (2012) refers to this as ‘system 1’ and ‘system 2’ processing, but the idea is the same: the brain does lots of complicated computations in the background, without us having to worry about it. Note that K&R’s BIG QUESTION only works to the extent that pronoun interpretation is a clear case of ‘system 2’ processing, in which interpreters need to plough their way through complicated, effortful computations. This is an empirical issue. In fact, the spirit of K&R’s probabilistic model leads me to think that K&R may well put most of the burden of pronoun interpretation on effortless ‘system 1’ processing. So, although the simplicity criterion is fine if reinterpreted as theoretical parsimony, what does not follow is that the processing itself also has to be simple. In terms of processing, things just need to be computable, by the brain, in time.

## 2 The method

Most of the evidence that K&R aim to explain with their model comes from passage completion. In the experiments, people are asked to complete fragments such as “Amanda amazed Brittany because she \_\_\_\_”, “John passed a comic to Bill. He \_\_\_\_”, or “John detested Bill. \_\_\_\_”, typically under the instruction “to imagine a natural story continuation for each prompt, writing the first continuation that came to mind and avoiding humor” (e.g., Rhode & Kehler, 2013). I have no principled objections against this method – fragment completion is one of the ways in which cognitive scientists can gain access to cognition. It is fast, cheap, and allows for web experiments, i.e., involving possibly lots of participants, and not just psychology and linguistics undergraduates. It is also representationally rich: where reading time, ERP and fMRI researchers struggle with a representationally empty signal, those that use fragment completion have contents at their hands. And, as pointed out by K&R, presumably more valid and reliable contents than the data generated by a single linguist exploring the hypothesis space behind a desk. The benefits are enormous.

Having said that, there’s also some risk involved in primarily using fragment completion data. First, and known as the “task demands” or “social desirability” issue in psychology, respondents may be inclined to give you what you want. Now, researchers will usually try to hide their desires, and they will often be successful. The socially minimal setting of a web experiment can certainly help here. Also, as illustrated by Rohde, Kehler & Elman (2007), actually *telling your participants* what you want – Tell us what happened next, or tell us why – can be part of the design. Still, in terms of spreading risks, it seems wise to complement a data type that is potentially very sensitive to task demands because of its content richness with data types that are representationally less rich (reading time, ERPs etc). Of course, the latter methods are not invulnerable to task demands either (see Van Berkum, 2004), but at least, the vulnerability is different.

There’s also a deeper concern associated with predominantly relying on passage completion data, which is that some phenomena may be overlooked, particularly those that involve the details of actual processing unfolding in time. An example comes from my own work on pronoun interpretation. As mentioned before, research with implicit causality verbs has shown that when people read “David praised Linda because . . .”, they rapidly anticipate more information about Linda, such that a subsequent masculine pronoun incurs immediate processing costs, revealed in reading behavior (Koornneef & Van Berkum, 2006) and brain potentials (Van Berkum et al., 2007). This is an instance of what K&R refer to as the NEXT-MENTION bias. However, recent EEG evidence (Van Berkum, De Goede, Van Alphen, Mulder, & Kerstholt, in press) reveals that the bias varies

with the participant's current mood: while readers in a good mood do show evidence for such referential anticipation, readers in a bad mood (as induced by depressing film fragments) no longer seem to anticipate who's going to be talked about next. Importantly, whereas the impact of mood on referential anticipation shows up in real-time reading assessed by EEG (experiment 1), it does *not* show up in passage completion (experiment 2). In the latter, something like "David praised Linda because . . ." is as biasing to readers in a bad mood as to those in a good mood.

Now, whether a *mood effect* on real-time referential anticipation is relevant to a theory of pronoun interpretation depends on the intended scope of the theory (see below). But for current purposes, the point is that potentially interesting modulations of referential anticipation and subsequent pronoun interpretation can remain totally invisible in passage completions. There is nothing principled about this: ERPs are not necessarily more sensitive than passage completion or other behavioral measures (see Van Berkum, 2004, for discussion). But in this case, a phenomenon that remains invisible in passage completion happens to show up in ERPs.

### 3 The ontology: Coherence only, or feelings too?

The mood findings obviously also raise another issue. In terms of the K&R theory: why would  $P(\textit{referent})$  depend on mood in one situation, but not in another? It is not that people in a bad mood shut down their entire cognitive system: subject-verb agreement violations elicit comparable ERP responses in both moods, and early ERP indices of attention are also unaffected by mood (Van Berkum et al., in press, experiment 1). In line with recent accounts of mood as a bio-energetic signal for available resources (e.g., Zadra & Clore, 2011), one viable explanation of why mood selectively modulates referential anticipation is that after having seen a negative-mood inducing film, people are simply less inclined to invest in exploratory, heuristic look-ahead. This account – which would incidentally imply that not all forms of pronoun-relevant anticipation have negligible costs – is consistent with recent behavioral evidence that a bad mood can attenuate the degree to which readers draw predictive inferences in text (Mirus & Beeman, 2012), and with N400 evidence that such a mood can also sometimes attenuate conceptual expectation in sentences and text (see Van Berkum et al., in press, for discussion and references). Regardless of the specific explanation, though, these mood findings reveal some of the unexpected complexity that can go into  $P(\textit{referent})$ .

This touches on a more general matter. Although the proactive K&R model is in several ways very different from classic reactive SMASH theories, it does seem

to share a wider ontology, which, as it happens, can also be traced back to big 1970's Cyber mainframes. For a long time, cognitive scientists have viewed the brain as a knowledge-level system (Newell, 1982), whose specialty was to compute accurate representations of the world. Translated to psycholinguistics, this computer-inspired perspective has led scientists to for example ask how comprehenders disambiguate word sense, construct a syntactic representation, compute reference, derive propositions or conversational implicatures, and update the situation model. It's all about building coherent models of the world.

However, the brain is not a scientist analyzing the environment in a dispassionate way (Damasio, 2010). Brains care about good and bad, about *valence*. And for good reason: staying away from bad stuff and approaching good stuff helps keep you alive, healthy, and attractive – and as such, it tends to help you generate offspring with the same useful inclinations. Special neural circuitry has evolved to handle valence in the brain, in a network that includes, amongst others, the amygdala, prefrontal cortex, the insula, cingulate cortex, the hypothalamus, nucleus accumbens, and the brainstem (see, e.g., Hamann, 2012; Panksepp & Biven, 2012). This 'emotional brain' or 'affective system' – the term 'limbic system' has gone out of fashion – is central to making you feel good or bad about certain things, and to as such guide your behavior in accordance with the goals set by biological evolution and/or you as a goal-directed individual. And such affect does *not* come as an afterthought. In cognitive science and neuroscience, it has become crystal clear that the affective system is involved in shaping and bending our mental representations right from the start (Barrett & Bar, 2009; Damasio, 2010; Haidt, 2012; Zadra & Clore, 2012). The brain doesn't actually aim for coherence in any dispassionate, objective sense, of the type that scientists generally aim for in their work. The brain is a deeply affective, subjective system.

The language sciences have long been at the heart of classic cognitive science. But with much of cognitive science turning increasingly affective, linguists and psycholinguists are beginning to wonder to what extent their field should reorient itself too (e.g., Foolen, 2012; Van Berkum, 2010; see Besnier, 1990, for a pioneering overview). Would an affective turn in the language sciences have any implications for theories of pronoun interpretation? I think it would. The mood study discussed before is a case in point. But there are other reasons why affect is likely to matter to pronoun interpretation. For one, language use is often social action (Clark, 1996; Levinson, 2006; Tomasello, 2008), and it is often also *about* social action (e.g., in gossip and fiction). And social action is not just coherent or not, it is saturated with affect. Outside of the language lab, knowing who did what to whom is not just about keeping track of X's and Y's in some disembodied DRT space – it is about keeping tabs on real people around you, or, in fiction, about models for what those real people might be like (Mar & Oatley, 2008). These

are people you like or dislike, love or fear, identify with or abhor. And to the extent that these X's and Y's are affectively loaded representations, they can draw additional attention (Barrett & Bar, 2009; Zadra & Clore, 2011), of the speaker and/or the hearer. That is, affective factors can influence  $P(\text{referent})$ . Furthermore, as indicated by honorifics and other, related phenomena, the very choice of referential expression itself can be a value-loaded social signal, allowing interlocutors to pay their respects and otherwise negotiate the relationship in the affective plane (Enfield, 2006; 2009). That is,  $P(\text{pronoun} | \text{referent})$  is not just controlled by information-theoretic factors such as topichood, but is also governed by social-affective factors.

My intention is not to bug K&R with “how could you forget about this?” comments – it takes courage to try to put together an explicit unifying theory, and if anything, such integrative attempts should be applauded. In fact, careful as they are, K&R themselves open the door to extensions: “Our analysis should not be construed as claiming that no other factors other than those we have discussed influence pronominal reference”. What I think should be avoided, though, is that all those extensions stay within the realm of discourse coherence. Yes, coherence is part of the story, we do try to build sensible situation models. But it's a big world out there, and if we are to understand how we use the language code to act in that world, we need to pull affective factors back into the equation too.

#### 4 The ontology: What happens next at what level?

There's a related point to be made. Going back to the Bill Murray conversation at the beginning, one can see that there are a zillion different things one can make predictions about. In principle, everything with a certain degree of systematicity can be predicted. However, in the context of language comprehension it may be helpful to ask whether particular predictions involve (a) developments in the domain under discussion, (b) references being made to those developments or the agents involved *by the communicator*, (c) the signs used by the communicator as part of this, and (d) the conversational and social ‘moves’ made by him or her. One can easily imagine partly coupled forward models running at each of these levels, particularly if they are implemented as ‘system 1’ operations, and largely rely on associative memory.

In text comprehension, the distinction between (a), (b) and (c) is a relatively clear-cut one. Imagine reading “When he found out the China museum had decided to sack him, the outraged employee grabbed their most precious and most delicate porcelain vase and violently hurled it towards a massive concrete wall”. The situation model constructed for this would involve lots of things, including a



nice fragile vase on its way to a museum wall. Of course, world knowledge allows us to predict that the vase will probably break (cf. Potts, Keenan & Golding, 1988), a prediction that essentially involves running a forward model of the (real or fictitious) world under discussion. However, such situation model predictions are different from predictions *about communication*, i.e., about what exactly the communicator will comment upon, and by means of which signs.

An example of the latter would be to predict a specific upcoming word, such as “broke” after continuing the story with “Of course, the vase . . .”; we know that people routinely do this (e.g., DeLong, Urbach & Kutas, 2005; Otten & Van Berkum, 2009; Van Berkum, Brown, Zwitserlood, Kooijman & Hagoort, 2005). This reveals something interesting about prediction in language comprehension, something that cannot easily be found in other areas of cognition, like vision. Because most instances of language use are referential, i.e., *about* something else, because language users can choose what to refer to and comment on, and because they can choose different signs to do the job, anticipation in language comprehension is typically multi-layered (see Altmann & Mirković, 2009, for a similar point). This is not to say that the underlying mechanisms must be different; perhaps they are not. But it is important to acknowledge that comprehenders usually deal with at least three different levels at which their targets of prediction exist: the domain under discussion, the particular aspect of that domain the communicator wishes to discuss, and the signs used to discuss it with.

As for (d), the conversational and social ‘moves’ made by the communicator, it is easier to return to the Bill Murray conversation. The anticipation of specific referential actions and signs is usually partly fueled by knowledge about plausible developments in the domain under discussion (typically referred to as ‘world knowledge’). But prediction in conversation involves something else too. For example, unless there are clear signals to suggest otherwise, we predict cooperative responses (Grice, 1975; Tomasello, 2008). We predict an answer after a question, and either acceptance or rejection after an invitation. Conversation is a collaborative project in which people *act* towards one another at every turn. Furthermore, conversation is not just a special type of social project, it also instantiates other social actions (think about having a violent debate, where people collaborate on communication to implement social distancing). Predicting those actions or ‘moves’ goes much deeper than merely predicting, say, an upcoming referent or referring expression; this is about what your interlocutor will *do*, not what her or she will point to or say exactly.

Why is all this relevant here? Again, the issue is what exactly goes into determining such terms as *P(referent)* or *P(pronoun | referent)* in the K&R model. In estimating the value of these ‘primitive’ terms for specific passages in completion

experiments, K&R can simply rely on the participants to do the math, somehow. But ideally, we'd also like to know *how* those participants arrive at their own estimates, via what other computations and using what other primitives. To work this out, it seems essential to develop an articulate view of the various different layers involved in anticipating upcoming communication. Presumably, such anticipation will involve *all* of the representational levels delineated by linguistics and psycholinguistics, plus a few more having to do with human sociality, human intentionality, and the emotions or preferences that guide and guard the two.

**Acknowledgment:** Supported by NWO Vici grant #277-89-001.

## References

- Altmann, Gerry T.M. & Mirkovic, Jelena (2009). Incrementality and prediction in human sentence processing. *Cognitive Science*, 33, 583–609.
- Bar, Moshe (2007). The proactive brain: Using analogies and associations to generate predictions. *Trends in Cognitive Sciences*, 11(7), 280–289.
- Bar, Moshe (2009) Predictions: A universal principle in the operation of the human brain (Introduction). Predictions in the brain: Using our past to generate a future (Moshe Bar Ed.) *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364, 1181–1182.
- Barrett, Lisa F. & Bar, Moshe (2009) See it with feeling: Affective predictions during object perception. Predictions in the brain: Using our past to generate a future (M. Bar Ed.) *Philosophical Transactions of the Royal Society B: Biological Sciences*, 364, 1325–1334.
- Besnier, Niko (1990). Language and affect. *Annual Review of Anthropology*, 19, 419–451.
- Clark, Herbert H. (1996). *Using language*. Cambridge: Cambridge University Press.
- Damasio, Antonio (2010). *Self comes to mind: Constructing the conscious brain*. New York: Pantheon.
- DeLong, Katherine A.; Urbach, Thomas P. & Kutas, Marta (2005). Probabilistic word pre-activation during language comprehension inferred from electrical brain activity. *Nature Neuroscience*, 8(8), 1117–21.
- Enfield, Nicholas J. (2006). *Social consequences of common ground*. In Nicholas J. Enfield, & Stephen C. Levinson (Eds.), *Roots of human sociality: Culture, cognition and interaction* (pp. 399–430). Oxford: Berg.
- Enfield, Nicholas J. (2009). Relationship thinking and human pragmatics. *Journal of Pragmatics*, 41, 60–78.
- Featherstone, Cara R. & Sturt, Patrick (2010). Because there was a cause for concern: an investigation into a word-specific prediction account of the implicit-causality effect. *The Quarterly Journal of Experimental Psychology*, 63(1), 3–15.
- Foolen, Ad (2012). *The relevance of emotions for language and linguistics*. In: Ad Foolen et al. (eds.), *Moving ourselves, moving others. Motion and emotion in intersubjectivity, consciousness and language*. Amsterdam: Benjamins, 349–368.
- Friston, Karl (2005). A theory of cortical responses. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1456), 815–836.

- Grice, Herbert P. (1975). "Logic and Conversation," *Syntax and Semantics*, vol. 3 edited by Peter Cole and Jerry L. Morgan, Academic Press. Reprinted as ch. 2 of Grice 1989, 22–40.
- Grosz, Barbara J.; Joshi, Aravind K. & Weinstein, Scott (1995). Centering: A framework for modeling the local coherence of discourse. *Computational Linguistics*, 21, 203–225.
- Haidt, Jonathan (2012). *The righteous mind: Why good people are divided by politics and religion*. London: Allen Lane.
- Hamann, Stephan (2012). Mapping discrete and dimensional emotions onto the brain: controversies and consensus. *Trends in Cognitive Neurosciences*, 16(9), 458–466.
- Hobbs, Jerry R. 1979. Coherence and coreference. *Cognitive Science*, 3, 67–90.
- Kahneman, Daniel (2011). *Thinking, fast and slow*. New York: Farrar, Straus & Giroux.
- Koornneef, Arnout W., & Van Berkum, Jos J. A. (2006). On the use of verb-based implicit causality in sentence comprehension: Evidence from self-paced reading and eye tracking. *Journal of Memory and Language*, 54(4), 445–465.
- Levinson, Stephen. C. (2006). *On the human "interaction engine"*. In Nicholas J. Enfield, & Stephen C. Levinson (Eds.), *Roots of human sociality: Culture, cognition and interaction* (pp. 39–69). Oxford: Berg.
- Mar, Raymond A. & Oatley, Keith (2008). The function of fiction is the abstraction and simulation of social experience. *Perspectives on Psychological Science*, 3, 173–192.
- Mesulam, Marsel (2008). Representation, inference, and transcendent encoding in neurocognitive networks of the human brain. *Annals of Neurology*, 64(4), 367–378.
- Metha, Mayank R. (2001). Neuronal dynamics of predictive coding. *Neuroscientist*, 7, 490–495.
- Mirous, Heather & Jung-Beeman, Mark (2012). *Stronger inference priming during story comprehension following positive mood induction than following anxious mood induction*. Submitted for publication.
- Newell, Allen (1982). The Knowledge Level. *Artificial Intelligence*, 18(1), 87–127.
- Otten, Marte & Van Berkum, Jos J.A. (2009). Does working memory capacity affect the ability to predict upcoming words in discourse? *Brain Research*, 1291, 92–101.
- Panksepp, Jaak & Biven, Lucy (2012). *The archaeology of mind: Neuroevolutionary origins of human emotion*. NY: Norton.
- Potts, George R., Keenan, Janice M. & Golding, Jonathan M. (1988). Assessing the occurrence of elaborative inferences: Lexical decision versus naming. *Journal of Memory and Language*, 27, 399–415.
- Pyykkönen, Pirta & Järvikivi, Juhani (2010). Activation and persistence of implicit causality information in spoken language comprehension. *Experimental Psychology*, 57(1), 5–16.
- Rao, Rajesh P.N. & Ballard, Dana H. (1999). Predictive coding in the visual cortex: a functional interpretation of some extra-classical receptive-field effects. *Nature Neuroscience*, 2(1), 79–78.
- Rohde, Hannah; Kehler, Andrew & Elman, Jeffrey (2007). *Pronoun interpretation as a side effect of discourse coherence*. In Proceedings of the 29th Annual Conference of the Cognitive Science Society. Nashville, TN, August 1–4, 2007.
- Rohde, Hannah & Kehler, Andrew (2013). *Grammatical and information-structural influences on pronoun production*. Submitted.
- Schacter, Daniel L.; Addis, Donna R. & Buckner, Randy L. (2007). Remembering the past to imagine the future: The prospective brain. *Nature Reviews Neuroscience*, 8(9), 657–661.
- Schooler, Lael J. & Anderson, John R. (1997). The role of process in the rational analysis of memory. *Cognitive Psychology*, 32(3), 219–250.

- Summerfield, Christopher; Egner, Tobias; Greene, Matthew S.; Koechlin, Etienne; Mangels, Jennifer A. & Hirsch, Joy (2006). Predictive codes for forthcoming perception in the frontal cortex. *Science*, 314, 1311–1314.
- Tomasello, Michael (2008). *Origins of Human Communication*. Cambridge, MA: MIT Press.
- Van Berkum, Jos J.A. (2004). *Sentence comprehension in a wider discourse: Can we use ERPs to keep track of things?* In Manuel Carreiras, Jr., & Charles Clifton (Eds.), *The on-line study of sentence comprehension: eyetracking, ERPs and beyond* (pp. 229–270). New York: Psychology Press.
- Van Berkum, Jos J.A. (2010). The brain is a prediction machine that cares about good and bad – Any implications for neuropragmatics? *Italian Journal of Linguistics*, 22, 181–208.
- Van Berkum, Jos J.A., Brown, Colin M., & Hagoort, Peter (1999). When does gender constrain parsing? Evidence from ERPs. *Journal of Psycholinguistic Research*, 28(5), 555–571.
- Van Berkum, Jos J.A.; Brown, Colin M.; Zwitterlood, Pienie; Kooijman, Valesca & Hagoort, Peter (2005). Anticipating upcoming words in discourse: Evidence from ERPs and reading times. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 31(3), 443–467.
- Van Berkum, Jos J. A.; De Goede, Dieuwke; Van Alphen, Petra; Mulder, Emma; Blokland, Yvonne & Kerstholt, José (in press). *How robust is the language architecture? The case of mood.* To appear in *Frontiers in Language Sciences*.
- Van Berkum, Jos J.A.; Koornneef, Arnout W.; Otten, Marte & Nieuwland, Mante S. (2007). Establishing reference in language comprehension: An electrophysiological perspective. *Brain Research*, 1146, 158–171.
- Zadra, Jonathan R., & Clore, Gerald L. (2011). Emotion and perception: The role of affective information. *Wiley Interdisciplinary Reviews: Cognitive Science*.