Anticipating the Garden Path:
The horse raced past the barn ate the cake

Gerry T.M. Altmann
Department of Psychology, University of York, UK.

> *In 1970, the Beatles disbanded, the Nuclear Non-Proliferation Treaty went into effect, the first jumbo jet went into service, the US became environmentally conscious (the Environmental Protection Agency was born), Simon and Garfunkel released* Bridge over Troubled Water*, Alexander Solzhenitsyn won the Nobel Prize for Literature, and Thomas G. Bever published* The Cognitive Basis for Linguistic Structures*. This was a lot to deal with for a ten-year old boy struggling with his maths homework. Shameful as it might seem (after all, the breakup of the Beatles* was *important), only one of these events had a traceable influence on that 10-year old's life. This chapter is an autobiographical account of that influence; it is a 1st person perspective on how the field of psycholinguistics, as seen by that developing 10-yr old, did itself develop.*

Cognitive Basis can be considered a backdrop against which much contemporary psycholinguistics can be viewed. Bever foresaw issues that would define a major part of the field (that part concerned with sentence processing) over the subsequent 40 years (and more). The following three quotations exemplify Bever's vision and contribution. And although the first is perhaps the better remembered, I shall argue over the course of this chapter that the reverse order is the more telling.

On the resolution of syntactic ambiguity:

> *The horse raced past the barn fell* (p. 316)

On constraint satisfaction:

> *… the most likely semantic organization among a group of phrases can guide the interpretation of sentences, independently of and in parallel with perceptual processing of the syntactic structure. […] In the actual application of language, specific contexts must provide far stronger immediate constraints and basis for prediction of the most likely meaning of a sentence independent of its form. Thus, most normal perceptual processing of sentences is probably carried out with little regard to actual sequence or structure; rather, the basic relational functions (actor-action- object-modifier) are assigned on the basis of temporary ("contingent") and generic ("constant") semantic probabilities."* (p. 297)

On the emergentist approach to language development:

> *Many aspects of adult language derive from the interaction of grammar with the child's processes of learning and using language. Certain ostensibly grammatical structures may develop out of other behavioural systems rather than being inherent in grammar. That is, linguistic structure is itself partially determined by the learning and behavioural processes that are involved in acquiring and implementing that structure"* (p. 280) *[…] the child may extract particular perceptual strategies by selective induction over his early linguistic experience* (p. 311)

To preview this chapter, I shall describe how all three of these observations have influenced my own perspective on the field of psycholinguistics, starting in the 1980s with syntactic ambiguity resolution. This early work considered processes at or beyond the point of ambiguity (e.g.

To appear in: M. Sanz, I. Laka, M. Tanenhaus (Eds.) *Language Down The Garden Path: The Cognitive and Biological Basis for Linguistic Structure*. Oxford: Oxford University Press.

Frazier, 1979; Frazier & Fodor, 1978; Fodor & Frazier, 1980; Rayner, Carlson, & Frazier, 1983). In terms of '*the horse raced past the barn fell*', the emphasis at this time was on how processes responsible for the interpretation of '*raced*' led to the subsequent incomprehensibility of '*fell*'. The emphasis changed, in the late 1990s, to consideration of processes occurring *before* the ambiguity, and how the interpretation of '*the horse*' had consequences for the subsequent interpretation of '*raced*' (Altmann & Steedman, 1988; Altmann, 1988; Altmann, Garnham, & Dennis, 1992; Crain & Steedman, 1985). Subsequently, there was a shift from studying *ambiguous* sentences to studying *unambiguous* sentences; sentences such as '*the boy will eat the cake*' (Altmann, 1999; Altmann & Kamide, 1999). The 1990s also saw the development of some of the most important insights that have shaped contemporary thinking on language processing (e.g. Elman, 1990; MacDonald, Pearlmutter, & Seidenberg, 1994). These emerged through considering Bever's third observation on the importance of learning (cf. Altmann & Mirkovic, 2009; Elman, 1990) and, from there, a fuller understanding of the origins of the constraints that form his second observation, on probability and prediction. The narrative that follows is a personal journey, signposted by *The Cognitive Basis of Linguistic Behavior*.

Background

In the 1970s and early 80s, the focus of much sentence processing work was on 'parsing preferences' and syntactic ambiguity resolution, stimulated by '*the horse raced past the barn*' and what became a large class of so-called "garden-path sentences". Bever had explained the garden path phenomenon in terms of a set of (perceptual) strategies that were experientially based (and which "*may reflect a statistical preponderance in actual utterances*", p.299). These strategies meant that the verb '*raced*' in the above example would be interpreted as a main verb (indicating that it was the horse that was doing the racing) rather than as a subordinate verb in a reduced relative clause construction (which is the correct interpretation, in fact; cf. '*the horse that was raced past the barn fell*', in which case the horse is raced, and someone else is doing the racing). However, the field took a different approach, proposing that these strategies reflected instead properties of the mental analogues of syntactic structure (cf. Kimball, 1973; Frazier, 1979). Subsequently, they were explained in terms of *architectural* divisions within the cognitive system; in effect, they were due to the structure of the memory systems in which language processing was grounded. Thus came into existence the Sausage Machine (Frazier & Fodor, 1978; Fodor & Frazier, 1980), and amongst the first computational linguistic models of sentence processing that attempted to explain psycholinguistic phenomena as an emergent property of the computational architecture in which analyses are pursued (Wanner 1980; 1987). Briefly, the Sausage Machine was a 2-stage parser consisting of a first stage that assigned lexical and phrasal nodes to the words it had access to in the sentence (through a limited-view window), and a second stage that took the 'packages' produced by the first stage and combined these to form a complete phrase marker. Wanner's models were extensions of a computational architecture developed in Wanner and Maratsos (1978) and based on Augmented Transition Networks (ATNs; Woods, 1970). Wanner and Maratsos (1978) proposed a computational architecture for describing how relative clauses could be processed (e.g. '*the horse that was raced past the barn fell*'). In an ATN there are "nodes" which reflect the internal state of the system, and arcs connecting these nodes which reflect the conditions that need to be met in order for the system to enter a particular state (corresponding, in effect, to a partial parse tree of the unfolding sentence). The ATN lends itself naturally to a physical metaphor in which a path through the network is traced as arcs are traversed, depending on which conditions are met and which sub-states are entered and subsequently left. By ordering the arcs in a way that was computationally maximally efficient (to simplify the argument somewhat), Wanner demonstrated that the preferences to interpret local syntactic ambiguities one way or another were a straightforward consequence of the ATNs architecture. And although the initial ATN model was serial (Wanner, 1980), a subsequent version implemented parallelism (Wanner, 1987). Both the Sausage Machine and ATN models of parsing preferences, and some of their shortcomings, are reviewed in more depth in Altmann (1988).

The focus on architectural explanations for parsing preferences was accompanied in the early 1980s by a perspective on cognition which assumed that cognitive functions are divided amongst a variety of informationally encapsulated sub-systems that process and then feed information to other sub-systems – J.A. Fodor's Modularity Hypothesis (Fodor, 1983). The Sausage Machine, for example, comprised informationally encapsulated modules that had no interaction with semantic, pragmatic, or discourse processes. A central claim at the time was that local syntactic ambiguities were initially resolved on the basis of various properties of the alternative syntactic structures that were possible, without appeal to higher-level semantic or pragmatic context. In the late 1980s and early 1990s, the field became dominated by a debate over whether syntactic ambiguities might nonetheless be resolved by appeal to such higher-level information. Crain and Steedman (1985) argued that such appeal did occur, and was computationally manageable, for a large class of ambiguities that had until then been explained through appeal solely to syntactic structure (Frazier, 1978; Kimball, 1973). Specifically, they argued that syntactic ambiguities are resolved through interaction with the context in which the ambiguous fragment occurs (cf. Bever's '*specific contexts must provide far stronger immediate constraints and basis for prediction of the most likely meaning of a sentence*'). Crain and Steedman (1985) and then Altmann and Steedman (1988; see also Altmann, 1988) argued that the *referential* context is critical in respect of one of the two major classes of garden path phenomena identified by Frazier (1979); namely, the class described by *Minimal Attachment* (the other class was described by *Late Closure*; cf. *Right Association*; Kimball, 1973).

Across a range of studies, we focused on ambiguities that arose through the interpretation of a noun phrase as either simple or complex – contrasting, for example, '*the girl watered the flowers with the silver watering can*' (a simple NP; the watering was done with the watering can) vs. '*the girl watered the flowers with the scented petals*' (a complex NP; the flowers had the petals). In the absence of any prior context, the complex NP case is interpreted initially as a simple NP structure, resulting in a garden path and elevated reading times on '*scented petals*' (because of the initial attempt to interpret this phrase as modifying the verb, rather than the preceding NP). We introduced a referential context in which there were two sets of flowers, one of which had scented petals, and showed that the garden path effect was eliminated (Altmann & Steedman, 1988; see also Altmann, 1988; Altmann et al. 1992). The essential idea here was that an interpretation of '*scented petals*' as an NP-modifier would resolve the referential ambiguity inherent in the phrase '*the flowers*' when that phrase was embedded in a context in which there was more than one set of flowers, and hence the preference to interpret '*scented petals*' as an NP-modifier, contrary to the usual preference to interpret it initially as a VP-modifier (cf. '*with the silver watering can*'). We thus showed that appropriate contextual (i.e. interpreted semantic) information could override the preferences that had hitherto been ascribed to an autonomous syntactic parsing module (see also Tyler & Marslen-Wilson, 1977, for perhaps the earliest demonstration of a contextual influence on parsing).

Bever had assumed that certain kinds of context could constrain sentence processing. But unlike Bever, we believed that semantic processing could not operate independently of syntactic processing (the second quotation from *Cognitive Basis* at the start of this chapter, on constraint satisfaction, clearly allows for such independence). We believed that the purpose of syntax was to tell you which semantic rules should be applied to which parts of the language (cf. Bresnan and Kaplan's "Strong Competence Hypothesis"; Bresnan, 1982), and that if semantics could operate independently of syntax, its acquisition would pose a considerable challenge to theories of language acquisition, not to mention theories of the relationship between syntax and semantics. In part, our adherence to this rule-to-rule relationship between syntax and semantics was borne from Steedman's development of Combinatory Categorial Grammar (CCG: Ades & Steedman, 1982; Ajdukiewicz, 1935; Steedman, 1996). CCG is a grammatical formalism which generates left-branching structures that are more amenable to incremental interpretation (and adherence to the Strong Competence Hypothesis) than were traditional grammars of the time. There are two other properties of CCG that made it particularly attractive at the time (and continue to make it attractive as a formal linguistic system): The first is that the lexical categories assigned to words can specify, in effect, the contexts in which those words can occur. The second (which is related to the first) is that the distinction between syntax and semantics is blurred, and the one-to-one relationship between 'syntax' and 'semantics' exists through the lexical categories assigned to words functioning both as syntactic and semantic categories (more or less – they reflect different levels of abstraction over the input); each category, and each combinatory function (which applied to those categories yields a combined structure), has both a syntactic and semantic interpretation.

A very brief aside: In contemporary psycholinguistics, the advent of connectionism at around this same time (the mid 1980s) was accompanied by the development of an 'emergentist' tradition in respect of language acquisition (see e.g. Elman et al., 1996; MacWhinney, 1999 – a tradition foreseen by Bever in the third quotation from *Cognitive Basis* at the start of this chapter, and discussed further below). In essence, this tradition assumes, in part, that the statistical patterning of words in the language (as well as other non-linguistic inputs to the cognitive system) underpins linguistic competence, and indeed, performance (the distinction between the two continues to vex psycholinguistics and linguists alike – early views on the distinction are reviewed in Valian, 1979). Within this tradition, the 'meaning' of a word reduces to a statistical analysis of the contexts in which that word can appear (cf. Burgess & Lund, 1997; Landauer & Dumais, 1997. These studies implemented the statistical approach via linguistic context alone, but the principles hold across non-linguistic contexts also). Of course, this approach to semantics is remarkably similar in some respects to approaches to syntactic category membership, in which membership of one category or another (e.g. noun or verb) is dependent on the contexts in which the word/category can occur. Thus (and now we come to the relevance of this aside), in the more 'contemporary' statistical approaches to language understanding, the lines between syntax and semantics are also blurred, and for similar reasons; they are each an abstraction over the input, and the contexts in which that input can occur (though not necessarily *the same* input – semantic abstraction will include non-linguistic context also). It was this blurred relationship between syntax and semantics that led us, at that time, to believe that semantics could not operate independently of syntax.

Paradoxically, despite the availability of a grammatical formalism in which the syntax and semantics worked almost as one, we shied away from an account of syntactic ambiguity resolution in which the two did actually work as one. Instead, we proposed an architecture in which syntax proposed alternative structural interpretations which semantics could adjudicate between on the basis of contextual fit. The corresponding mantra (from Crain & Steedman, 1985) was '*syntax proposes, semantics disposes*'. Crucially, we believed at the time that semantics could not tell the syntactic parser which analysis to pursue in the first place, although it could cause the parser to abandon a hypothesis: Thus, at '*the flowers*' in '*the girl watered the flowers…*', syntax could offer up either the simple or complex NP analysis of '*the flowers*', and semantics could, on the basis of the referential context and a mechanism for incremental referential constraint satisfaction, determine which analysis should go forward, and which should be terminated. But in a sense we 'got it wrong'; we believed that the theoretically interesting 'action', so to speak, happened either at or after the point of ambiguity. At the time, we didn't stop to think that the really interesting action might happen beforehand. And not just beforehand in the sentence (e.g. at '*watered*'), or in the text/discourse (e.g. when one or more lots of flowers were introduced). But beforehand in the experience (aka life) of the language user.

Adult language processing as an emergent property of the child's learning

At the end of the 1980s, Jeff Elman described a computational model of learning that revolutionized the way many people conceived of the relationship between learning, language acquisition, syntactic structure, and meaning. In so doing, Elman linked issues that pervade the three quotations that started this chapter. Elman (1990) modified the architecture of the simple recurrent network (SRN) as originally developed by Jordan (1986), and devised a novel task for teaching the network about language. These developments have been described extensively elsewhere (see Altmann, 1997, for a non-specialist account). Briefly, the SRN is a connectionist network in which autonomous units receive a signal and pass that signal to units at the next level 'up' in the network. The signal is modulated by the weight of the individual connection through which it passes. In its simplest form, the SRN consists of a layer of input units, a layer of output units, and a 'hidden' layer that mediates between the two. Critically, the hidden layer receives input not only from the input layer, but from a copy of its own activity at the previous time-step. Thus the activity at the input layer is determined not just by what it receives from the input at that moment in time, but by a memory of its own activity from the moment before. The network thus encodes ‘echoes’ of its past internal states. Teaching the network was simple: the network had to predict, given its input, what the *next* input would likely be (and a version of back-propagation, a standard learning algorithm in connectionist networks, would modify the weights on the internal connections so as to progressively reduce the error between the network’s predictions and the actual next inputs). So for a sentence such as ‘*the horse ate the hay*’, such a network would, if 100% successful, predict after ‘*the*’ the words ‘*horse*’ or ‘*hay*’, and after ‘*horse*’ ‘*ate*’, and so on. Of course, it could never learn with 100% accuracy, because the same word might be followed on different occasions by different other words. But what the network *could* learn was the range of words that could follow any particular input. And as mentioned earlier, knowledge of the range of words that can appear in any particular position within a sentence (i.e. in a particular context) constitutes knowledge of the syntactic category of that word (e.g. nouns occur in certain contexts, and verbs in certain others). But that knowledge also conveys semantic information – only certain words could occur after ‘*ate the*’; constituting in this case the semantic class of edible things (cf. the earlier discussion of the potentially blurred nature of syntactic and semantic categories). Elman’s implementation of the SRN was thus critical in driving the emergentist approach to language acquisition (Elman et al., 1996), as it provided an important demonstration of how structure could ‘emerge’ through abstraction across experience. It was also a part of a larger shift that led to the *constraint satisfaction* approach to sentence processing (MacDonald, Pearlmutter, & Seidenberg, 1994; Trueswell & Tanenhaus, 1994), in which sentence processing consists of the application of probabilistic constraints, in parallel, as a sentence unfolds, with no single constraint being more or less privileged than any other except in respect of its probabilistic strength.

But the relevance of Elman’s work here was that it showed that there existed a computational mechanism through which contextual influences on the unfolding interpretation of a sentence were tractable and essentially deterministic. The mechanism for ‘interpreting’ an unambiguous string of words was no different from that for interpreting a string that contained a local ambiguity. On encountering, for example, ‘*boys eat*’, the SRN would activate at its output layer representations corresponding to all the words that in its experience might follow the verb eat (in the context also of contexts in which boys were doing things). Similarly, On encountering, for example, ‘*horses raced*’, it would do the same, and the different continuations would reflect the contextual prevalence of one kind of continuation versus another. Ambiguity, in this model, is manifest as the space of possible continuations from a given point in a sentence. There is, therefore, little difference between an unambiguous fragment and an ambiguous one; the space of continuations is determined by the network’s past experience, as well as, of course, the current context. The power of the model lies in its predictive (and parallel) activation of what can come next.

The principles underlying the model’s workings led us, in Altmann and Mirković (2009), to conceptualize sentence comprehension in the following way:

> *‘Knowledge’’ of the language can be operationalized as the ability to predict on the basis of the current and prior context (both linguistic and, if available, nonlinguistic) how the language may unfold subsequently, and what concomitant changes in real-world states are entailed by the event structures described by that unfolding language. Such predictions constitute the realization of the mapping between sentence structures and event structures.*
>
> *Concurrent linguistic and nonlinguistic inputs, and the prior internal states of the system (together comprising the context), each ‘‘drive’’ the predictive process, and none is more privileged than the other except insofar as one may be more predictive than the other with respect to the subsequent unfolding of the input.*
>
> *The representation of prior internal states enables the predictive process to operate across multiple time frames and multiple levels of representational abstraction. The ‘‘grain size’’ of prediction is thus variable, with respect to both its temporal resolution and the level of representational abstraction at which predictions are made.*
>
> Altmann and Mirković (2009), p. 586.

The SRN opened up the possibility of an alternative route to investigating the role of context in sentence processing – not just in terms of how context might be used to resolve syntactic ambiguity, but in terms of how it might be used to drive prediction during the processing of *unambiguous* sentences. So hearing '*the boy will eat*' should activate the prediction that, most likely, a noun phrase will follow, and that it will refer to something that was edible. Similarly, hearing '*the toddler will drink*' should activate the prediction that, most likely, a noun phrase will follow that will refer to something that can be drunk, and moreover, something that would most plausibly be drunk by a toddler. But is this what happens?

Prediction during sentence comprehension: Empirical findings

Much of the work on prediction during sentence comprehension has been based on either the 'Visual World Paradigm' (Cooper, 1974; Tanenhaus et al. 1995) or on studies measuring event-related brain potentials (ERP). The findings can be summarised briefly as follows: In the context of a visual scene, hearing a sequence such as '*the woman will drink*…' causes the eyes to move, at '*drink*', to whatever in that scene is drinkable, such as a glass of wine (Altmann & Kamide, 1999). In fact, even if the scene is removed *before* the spoken sentence begins to unfold, exactly the same pattern is found, with eye movements at '*drink*' towards where the wine *had been* (Altmann, 2004). Given that language often refers to things that are not concurrent with the unfolding language, this last result is an important demonstration of how the paradigm generalizes beyond concurrent reference. These studies suggest that in appropriate contexts, we anticipate at the verb what might follow. In a reading study that was equivalent in some respects to these two eye tracking studies (Altmann, 1999), participants read '*The woman looked at the wine. She drank*…' or '*the woman looked at the sign. She drank…*'. Reading times were longer at '*drank*' in the case where nothing drinkable had been mentioned in the context (the sentence continued '*…some wine that she'd brought with her*'). In this case, it appeared that much like in the eye tracking studies, participants assumed that the object of the verb (i.e. the thing that would be drunk) would be drawn from the context, and if there was nothing appropriate in the context, reading times were elevated (and there were increased 'stops making sense' judgments). Exactly why it is assumed that the anticipated objects will be drawn from the context (visual or discourse) has been explained elsewhere (Altmann & Mirković, 2009).

These effects are not simply due to something about verbs: First, for a scene showing a glass of wine, a glass of milk, a woman and a toddler, hearing '*The toddler will drink the milk*' engenders looks to the milk at '*drink*', whereas hearing '*The woman will drink…*' engenders looks to the wine. Thus the eye movements reflect here the *combination* of the verb with its prior subject; they reflect whatever will plausibly be drunk *by the person doing the drinking*. Second, sequences of nouns in Japanese (a language in which the verb appears at the end of each sentence) have been show to activate predictions about the kind of verb that can follow (Kamide et al., 2003), and the Altmann (1999) result has been reinterpreted as the context in the first sentence ('*The woman looked at the wine*') restricting the range of actions or events that could be referred to next – in effect, predicting the range of verbs that could come up next (Altmann & Mirkovic, 2009; see also McRae, Hare, Elman, & Ferretti, 2005; who showed that typical participants in events prime the verbs that denote the event action).

These data appear straightforward, at least empirically so – you hear something, and look at something else that's likely to be referred to next, and this is deemed to be evidence of anticipatory processing. But before addressing *what* is anticipated, it is worth stopping for a moment to consider whether such data necessarily reflect anticipation. We showed, in Altmann & Kamide (1999) that hearing *'eat*' in '*The boy will eat the cake*' engenders looks to a piece of cake. But Yee & Sedivy (2005) showed that hearing '*bread*' would *also* engender looks to a piece of cake. Accounts of *why* the eyes move to cakes on hearing '*eat*' or '*bread*' (Altmann & Kamide, 2007) do not distinguish between one case and the other – they are both due to overlap between the conceptual representation activated by the spoken word(s) and the conceptual representation previously activated by the picture of the cake.[1] Altmann and Mirković (2009) equated the change in state that occurs when the conceptual representation of '*eat*' meets the conceptual representation of cake as 'thematic role assignment' (with the cake being anticipated to be the object of the eating). But few would also equate the change in state that occurs when the conceptual representation of '*bread*' meets the conceptual representation of cake as *thematic role assignment*; instead, it would be equated with the overlap in conceptual structure between cakes and bread (which, in other paradigms, is equated with semantic priming). However, the fact that behaviourally they may come about through the same mechanism raises the question of whether the difference may be a difference in nomenclature alone. Thus, as we saw in the earlier discussion of lexicalist approaches to grammar (of which CCG is one example), the boundaries between lexical and sentential processes are yet again blurred. But the *eat/bread*-cake facts raise the issue of the criteria we should adopt for labelling a representation as 'anticipatory'. Is this also just a matter of nomenclature?

The answer to this last question is '*no*'. The conceptual overlap in the representations elicited by the cake and by the word '*bread*' is due to the commonality they share in respect of the contexts in which they have been experienced – activating one necessarily activates the overlapping components of the other. However, the overlap in the representations elicited by the cake and by the words leading up to and including '*eat*' is experientially different: The meaning of a verb is composed, in part, from the meanings of the nouns it co-occurs with; the meaning of a noun is composed, in part, from the meanings of the verbs it co-occurs with (just as the meaning of an object is composed in part from the nature of the events and actions it can engage or be engaged in – i.e. its *affordances*). Thus, the conceptual overlap between '*eat*' and '*cake*' is due to the abstract representational properties of cake that, via abstraction across the contexts in which cakes (and '*cake*') can be experienced, encode aspects of eating. And although conceptual overlap entails activation of one by the other, such that seeing a cake and hearing '*the boy will eat*' causes a boost in activation for the representation corresponding to the cake (see above), and seeing a cake and hearing '*bread*' causes a boost in activation for that same representation, these boosts in activation in the two situations reflects different things: In the first case, it reflects a representation of the concomitant changes in real-world states that would correspond to the unfolding event described by the language; it is a prediction of what may come/happen next. In the second case, it reflects the alternative possible worlds in which, in effect, the cake could *replace* the bread. Although the mechanism of activation is the same, the information that the activation reflects is distinct. Of course, one can ask how the system (whether the human cognitive system or a bespoke SRN) can distinguish one from the other. But the '*eat*'-cake dependency exists by virtue of a trajectory through time and state-space. The '*bread*'-cake dependency exists by virtue of the shape of the current state-space (its 'surface') – they are thus quite distinct. Nonetheless, as in the case of previous distinctions we have considered, the boundaries are, once again, a little blurry. Indeed, Kukona, Fang, Aicher, Chen, and Magnuson (2011) report evidence which suggests that in cases like '*eat the cake*' there is likely to be both active prediction but also non-predictive priming from the verb to the noun.

If we take as given that the current state of the system constitutes a prediction of which states the system may enter in the future (and hence, what corresponding input would be required to put it in those states), we can ask what it is that is being predicted (or anticipated; see Altmann & Mirkovic, 2009, for a theoretical distinction between the two – here, the two are used interchangeably).

[1] This is not to say that looks to the cake on hearing '*eat*' are purely lexically driven (i.e. due *only* to the conceptual overlap between '*eat*' and cakes) – the fact that anticipatory eye movements are driven by combinations of words (noun-verb, Altmann & Kamide, 1999; noun-noun, Kamide et al., 2003) or are modulated by the tense of the verb ('..*will drink..*' engendering looks to a full glass, but '..*has drunk..*' to an empty glass; Altmann & Kamide, 2007) suggests that these effects are not purely lexical.

Prediction during sentence comprehension: What is being predicted?

In principle, the anticipatory eye movement data reported above could reflect predictions of the upcoming language – representations of the form of the words that will likely be heard next – or they could reflect predictions about the unfolding *conceptual correlates* of the event which that language describes – representations of the concomitant changes in the real-world that would constitute the event described by the language. Or, the eye movement data could reflect both.

Perhaps the most compelling evidence that the (phonological) form of upcoming words can be anticipated comes from two similar, but independent, studies published by van Berkum, Brown, Zwitserlood, Kooijman, and Hagoort (2005) using ERP and spoken language, and by DeLong, Urbach, and Kutas (2005) using ERP and written language. To briefly summarize the latter study, they showed that for a sentence fragment such as '*The day was breezy so the boy went outside to fly…*' readers anticipate that a following article is more likely to be '*a*' than '*an*' – '*a kite*' is a more plausible continuation than '*an airplane*', and they found that the greater the likelihood, across their various stimuli, of continuations like '*airplane*' (i.e. requiring '*an*'), the smaller the N400 component at the preceding article (i.e. the smaller the mismatch between their expectations and what they in fact subsequently read).

Evidence that *conceptual* structure associated with the event itself can be anticipated comes from a series of eye movement studies described in Altmann & Kamide (2009). In one of the experiments, we showed participants scenes depicting amongst other things, a table, a bottle of wine and a glass both on the floor, and a woman. The scene was removed before participants would hear either '*The woman will move the glass onto the table. Then…*' or '*The woman is too lazy to move the glass onto the table. Instead…*' After one or other of these sentences, they heard '*she will pick up the bottle and pour the wine carefully into the glass*'. Both during the sentence-final *glass* in the second sentence, and beforehand as they *anticipated* the location of the pouring, the eyes moved towards where the table had been if that is where the glass was heard to have moved, or towards where the glass had actually been if the glass had not been described as moving. The anticipatory eye movements towards the event-specific location of the glass, rather than towards the actual location of the glass as represented in visual memory, suggest that the conceptual correlates of the event described by the unfolding language can indeed be anticipated.

Taken together, these data illustrate what in Altmann and Mirković (2009) we described as predictive processes operating across multiple time frames and multiple levels of representational abstraction. The Altmann and Kamide (2009) data directly demonstrate our operationalization of competence in a language as being the ability to predict the concomitant changes in real-world states that are entailed by the event structures described by that unfolding language. But with this characterization of competence in a language comes a hitherto ignored problem: Predicting the possible changes in real-world states entails representing those changes. But this poses a significant challenge, both to human cognition and to computational models of such cognition.

Event comprehension and the challenge of change

Many, perhaps even most, events entail change. But how would one go about representing change in, for example, an SRN? Prediction or anticipation in an SRN is reflected in activation patterns across the hidden layers at time *t*, contingent on inputs at time *t-1* (and earlier), which enable activation patterns across the output at time *t* (and beyond) that reflect time *t+1* (and beyond). But the hidden layer is a homogenous substrate, with activation patterns at time *t* laid down *on top of* the activation patterns at time *t-1*. The human brain may or may not do something similar – that is, 'overwrite' activity at one moment in time with activity at the next. Whether it does or does not, there is still an issue with respect to how we keep the distinct patterns distinct. To put this in concrete terms, consider the stimuli from a study by Hindy, Altmann, Kalenik, & Thompson-Schill (under review). These stimuli were in fact motivated by the *move-the-glass* stimuli from Altmann and Kamide (2009) but avoid confounding change in state with change of location:

The squirrel will crack the acorn, and then it will lick the acorn.
The squirrel will crack the acorn, but first it will lick the acorn.

How do we keep apart the representation of the acorn *before* it is cracked, and the representation of the acorn *after* it is cracked? Both are available for subsequent reference (i.e. at the second mention of ‘*the acorn*’). An overlapping substrate inevitably entails *competition* between one representation and another – depending on the circumstances, one representation must be selected at the expense of the other; in effect, the representations ‘compete’ for ‘attention’ (cf. the view of attention and competition described in Cohen, Aston-Jones, & Gilzenrat, 2004). But more than this, an overlapping substrate suggests that there may be *interference* between one representation and the other (to use an earlier example: Partial activation of the conceptual representation of cake may in certain circumstances impede the intended activation of the conceptual representation of bread). In the language and memory literature, similarity-based interference is well-attested (Bower, 2000; Gordon, Hendrick, & Johnson, 2001; Postman, 1971; van Dyke & McElree, 2006). If an object has to be represented in the ‘before’ and ‘after’ (and the minimum representation of an event entails such representation), these distinct representations will be inherently similar. Conceivably, then, representing a change of state may be more ‘costly’ than not having to represent such a change; to go back to the acorn example: There may be a cost to representing the before and after states of the acorn, making the acorn less accessible than if the event had left it unchanged.

The claim here is that event representations require the representation of multiple instantiations of the same object, each representing a distinct event-specific state of the object that underwent change as a consequence of the event. That these distinct instantiations might interfere with one another is reminiscent of proposals by Zaichik (1990) and e.g. Radvansky, Krawietz, and Tamplin (2011), who suggest that there can be competition and conflict between distinct event models (corresponding to e.g. the cracking event and the licking event). Here, however, we propose that interference may obtain not between event models/representations *per se*, but specifically between object representations, insofar as distinct representations of the alternative states of the same object will compete.

If distinct states of the same object must be represented, how might these states in fact be represented? The answer to this question has, as we shall see, empirical consequences. There are (at least) three broad possibilities: First, the cognitive system maintains multiple time-stamped ‘copies’ of the object’s representation, each reflecting its state at that time, each bound to some subset of the information associated with the event as a whole, and each bound to the other along various dimensions (e.g. time, causation, space; see Zwaan & Radvansky, 1998, for the role of such dimensions in event representation). Second, the system maintains only a single object representation, but its featural composition (cf. featural approaches to semantic cognition; McRae, de Sa, & Seidenberg, 1997; Rogers & McClelland, 2004; Tyler & Moss, 2001) contains components that are themselves time-stamped, with different featural properties of the object representation reflecting different states at different times. In many respects, this is representationally equivalent to the first possibility, requiring that the different time-stamped components are bound to each other and to other dimensions of the event. Whereas the first possibility entails competition, and perhaps interference, between entire object representations, the second possibility entails competition, and perhaps interference, between the components of individual object representations. For the purposes of present discussion, we describe both these options as reflecting “multiple instantiations” of (all or parts of) the same object.

The third possibility is that in fact there are no multiple instantiations; there is just one. On this view, the cognitive system computes the contextually appropriate representation as necessary, using the initial object representation in conjunction with the event representations to which that object must (as in the multiple instantiation accounts) be bound. For the first of the *acorn* examples above, the system might represent just the cracked acorn, after hearing or reading ‘*the squirrel will crack the acorn*’, but if required, it would “transform” this into a representation of an uncracked version of the acorn. This would be required, for example, at the end of ‘*but first it will lick the acorn*’. In terms of theories of *simulation* (e.g. Barsalou, Simmons, Barbey, & Wilson, 2003), the simulation would in effect be reversed. If the continuation were instead ‘*and then it will lick the acorn*’, no transformation of the representation of the cracked acorn would be required. Conceivably, the system might instead represent, at the offset of ‘*the squirrel will crack the acorn*’, only the uncracked version, transforming it into the cracked version only when required (e.g.

when having to retrieve the representation at the end of '*and then it will crack the acorn*'). In this case, no transformation would be required in the '*but first*' case.

The first two possibilities predict that in both the acorn cases above (the '*and then*' and '*but first*' versions), there could be competition and interference between the distinct object instantiations. The third possibility – the single instantiation account – predicts that any increased processing load (due to the transformation of that instantiation into a new version) would interact with the temporal order of the cracking and the licking. So which is it? What *is* the empirical evidence?

The challenge of change: Empirical findings

In the Hindy et al. (under review) study, we adopted the following strategy to answer this question: A task that is known to engender competition, interference, and *conflict* is the Stroop task (Stroop, 1935). In this task, participants read a word such as "green" and have to respond to the color (by pressing a button) in which the letters are printed. When the colour is in conflict with the word (e.g. "green" is in red lettering), reaction times are slower than when the colour is congruent with the word ("green" is in green lettering). Much is known about this task (see Macleod, 1991). Recently, it has become increasingly used in identifying brain areas implicated in resolving conflict during the selection of semantic alternatives during sentence processing (e.g. January, Trueswell, & Thompson-Schill, 2009). Consequently, we asked whether the same brain areas that are sensitive to Stroop conflict would be sensitive to the predicted interference in the acorn examples above. We contrasted the following pairs of sentences:

The squirrel will crack the acorn, and then it will lick the acorn [acorn undergoes substantial change]
The squirrel will sniff the acorn, and then it will lick the acorn. [acorn undergoes minimal change
The squirrel will crack the acorn, but first it will lick the acorn. [acorn undergoes substantial change]
The squirrel will sniff the acorn, but first it will lick the acorn. [acorn undergoes minimal change]

We predicted that the *same* voxels that are sensitive to Stroop conflict would also be sensitive to the difference between the '*crack*' and '*sniff*' versions above *if* multiple instantiations of the acorn are indeed subject to mutual interference. Under the single instantiation account, we would predict a difference between the '*and then*' and '*but first*' cases, with any differences between the '*crack*' and '*sniff*' cases reflecting the cost of computing the new representations as required. As expected, we *did* find that the same voxels that were sensitive to Stroop conflict (in left posterior ventrolateral prefrontal cortex) were sensitive to the '*crack*'/'*sniff*' alternation. Moreover, there was *no* interaction with temporal order (the '*and then*' or '*but first*' versions), suggesting that the single-instantiation account is unlikely. We concluded that multiple instantiations of (all or parts of) the same object *are* represented (one way or another – see the two possibilities above), and that there is most likely interference between the instantiations. In fact, we found something else too: the amount of change that the object underwent, as found in offline ratings, predicted the amplitude of the BOLD response in those Stroop-sensitive voxels; the greater the change, the greater the BOLD response. In the Stroop task, a larger BOLD response is associated with greater conflict (and greater interference between the alternative representations). Thus, the correlations that we observed between degree of change and BOLD response suggest that the greater the change that the object underwent, the greater the conflict. It was not *similarity*-based interference that we found, but *dissimilarity*-based interference. Presumably, because the greater the change, the more semantic dimensions in conflict between the instantiations.

We found exactly the same pattern of results with stimuli such as:

The girl will stamp on the egg. Then, she will look down at the egg. [substantial change]
The girl will stamp on the penny. Then, she will look down at the penny. [minimal change]

Once again*,* the BOLD response in Stroop-sensitive voxels was predicted by the degree of change that the critical objects underwent – eggs tend to undergo more change when stamped upon than pennies. This second result shows that the effects observed in the '*crack*'/'*sniff*' cases were not due simply to differences between the verbs '*crack*' and '*sniff*' (or their equivalents across the experimental stimuli).

The empirical evidence suggests that event representation, and specifically the representation of change, does indeed come at a cost: The need to maintain multiple representational instantiations of (all or parts of) the same object engenders representational conflict. Although in these first studies we used fMRI to probe such conflict, ongoing behavioral studies with the visual world paradigm and with self-paced reading confirm a cost introduced by changing the states of objects that subsequently must be retrieved. However, if it is indeed the case that we construct multiple representational instantiations of the same objects, we are no closer to understanding how these are represented within the cognitive system; how might a system which over-writes itself at each moment in time keep such instantiations distinct? How could it *retrieve* the appropriate instantiations from within the echoes of its past states? How are the distinct instantiations bound to one another through whatever episodic event-specific knowledge might be available? These are not questions about language processing, but about the *conceptual correlates* of the world that we are able to describe through language. This is no longer psycholinguistics. It is *cognition.*

Where does this leave us? We've gone from horses racing past barns to squirrels cracking acorns. We've gone from syntactic ambiguity, constraint satisfaction, and emergentist theories of language development, to event representation and the challenge of change. To what extent do these shifts in topic across the past 40 years indicate actual *progress* towards understanding the cognitive basis for linguistic structure? And where might this lead us next?

From horses to squirrels: from predicting human behaviour to understanding the human mind.

The Cognitive Basis for Linguistic Structure was, 40 years ago, the intellectual equivalent of a starting pistol. The horse raced, and so did much of the field. Many careers were launched by the echoes of that single pistol shot. The field gradually transformed from being preoccupied with different kinds of ambiguity to being preoccupied with the different constraints that might help resolve such ambiguities in real time. How such constraints might be *implemented*, within computational systems, became an equally important issue as *Parallel Distributed Processing*, in its various guises, swept through the field. Developments in computational linguistics (*computational cognition* might be a better name for much of it), and particularly through a clearer understanding of how *development* might proceed within such computational systems, led to major insights concerning the emergence of representation as an abstraction across experience. The field has largely accepted the probabilistic constraints satisfaction approach to sentence processing, although their computational implementation remains debated (cf. Connectionist vs. Bayesian approaches to cognition – see the special issue of *Trends in Cognitive Sciences* on *Approaches to Cognitive Modeling, 2010,* ***14***). Where progress remains to be made is in considering, for example, how event representations are constructed incrementally as each sentence in the language unfolds, and how these representations build upon prior experience. The "challenge of change" identified above is just one example puzzle of the many that remain unsolved. It is one thing to demonstrate that there is a phenomenon worth exploring (e.g. showing that multiple representational instantiations of the same object do compete), but it is quite another to explain the cognitive underpinnings of such phenomena.

Forty years ago, Bever had a vision. That vision has largely been borne out by how the field has linked together, like he did in a single chapter, concepts from a diverse set of research traditions. Consideration of syntactic ambiguity led to consideration of multiple interacting constraints, of the relationship between syntax and semantics, and of the relationship between these two and issues in development. Syntactic ambiguity was, for many, the starting point. But syntactic ambiguity was a *tool* to probe the architecture of cognition. New tools have emerged but the puzzles remain the same. We *do* have a greater understanding of the emergent nature of cognitive representation and the processes that transform such representations moment-by-moment as the language we listen to, and the world we look to, unfolds before us. But our knowledge is limited. Progress occurs through the discovery of our limitations. And our continuing discovery of the many puzzles that must still be solved.

Forty years ago, a 10-yr old boy could not imagine the impact that one Thomas G. Bever would have on his and many other lives. That boy's early intellectual beginnings, embedded then in syntactic parsing, have led now to consideration of issues that impact on cognition more generally, and on the manner in which we do something that is surely even more fundamental than our ability to use language; to represent the changing world around us.

Acknowledgements

The research described in this chapter has been funded at different times by the Medical Research Council, The Economic and Social Research Council, and the Wellcome Trust. None of the work would have been possible without the input from collaborators, colleagues, and friends that kept this work alive over the past 30 years. They know who they are. I have been privileged to work with them. Just as I have been privileged to be a part of the legacy that is *The Cognitive Basis for Linguistic Structure.*

References

Ades, A., & Steedman, M. J. (1982). On the order of words. *Linguistics and Philosophy*, **4**, 517-528.

Ajdukiewicz, K. (1935) Die syntaktische Konnexitat. *Studia Philosophica*, 1, 1-27. English translation in Storrs McCall (ed) *Polish Logic: 1920-1939*, pp207-231, Oxford University Press, Oxford.

Altmann, G. T. M. (1997). *The ascent of Babel: An exploration of language, mind, and understanding*. Oxford: Oxford University Press.

Altmann, G. T. M. (1999). Thematic role assignment in context. *Journal of Memory and Language*, **41**, 124-145.

Altmann, G. T. M. (2004) Language-mediated eye movements in the absence of a visual world: The 'blank screen paradigm'. *Cognition.* **93**, 79–87.

Altmann, G. T. M., & Kamide, Y. (1999). Incremental interpretation at verbs: Restricting the domain of subsequent reference. *Cognition*, **73**(3), 247-264.

Altmann, G. T. M., & Kamide, Y. (2007). The real-time mediation of visual attention by language and world knowledge: Linking anticipatory (and other) eye movements to linguistic processing. *Journal of Memory and Language,* **57**, 502-518.

Altmann, G. T. M., & Steedman, M. J. (1988). Interaction with context during human sentence processing. *Cognition*, **30**(3), 191-238.

Altmann, G. T. M., Garnham, A., & Dennis, Y. (1992). Avoiding the garden path: Eye movements in context. *Journal of Memory and Language, **31**, 685-712.*

Altmann, G.T.M. (1988). Ambiguity, parsing strategies, and computational models. *Language and Cognitive Processes*. **3**(2), 73-97.

Altmann, G.T.M. and Kamide, Y. (2009). Discourse-mediation of the mapping between language and the visual world: eye-movements and mental representation. *Cognition*, **111**, 55-71.

Altmann, G.T.M. and Mirkovic, J. (2009). Incrementality and prediction in human sentence processing. *Cognitive Science*, **33**, 583-609.

Altmann, G.T.M., Garnham, A., & Dennis, Y. (1992). Avoiding the garden path: Eye movements in context. *Journal of Memory and Language*, **31**, 685-712.

Barsalou, L. W., Simmons, W. K., Barbey, A. K., & Wilson, C. D. (2003). Grounding conceptual knowledge in modality-specific systems. *Trends in Cognitive Sciences,* **7***,* 84–91.

Bever, T. G. (1970). The cognitive basis for linguistic structures. In J. R. Hayes (Eds.), *Cognition*

and the Development of Language. New York: John Wiley and Sons.
Bower, G.H. (2000). A brief history of memory research. In E. Tulving & F.I.M. Craik. (Eds.), *The Oxford Handbook of Memory*, pp. 3-32. Oxford University Press.
Bresnan, J. (Ed.). (1982). *The Mental Representation of Grammatical Relations*. Cambridge, MA: MIT Press.
Burgess, C., & Lund, K. (1997). Modelling parsing constraints with high-dimensional context space. *Language and Cognitive Processes*, **12**, 177-210.
Cohen, J. D., Aston-Jones, G., & Gilzenrat, M. S. (2004). A systems-level perspective on attention and cognitive control: Guided activation, adaptive gating, conflict monitoring, and exploitation vs. exploration. In M. I. Posner (Ed.), *Cognitive neuroscience of attention* (pp. 71–90). New York: Guilford Press.
Cooper, R. M. (1974). The control of eye fixation by the meaning of spoken language: A new methodology for the real-time investigation of speech perception, memory, and language processing. *Cognitive Psychology,* **6**, 84-107.
Crain, S., & Steedman, M. J. (1985). On not being led up the garden path: the use of context by the psychological parser. In D. Dowty, L. Karttunen, & A. Zwicky (Eds.), *Natural Language Parsing: Psychological, Computational, and Theoretical perspectives*. (pp. 320-358) Cambridge: Cambridge University Press
DeLong, K.A., Urbach, T.P., & Kutas, M. (2005). Probabilistic word pre-activation during language comprehension inferred from electrical brain activity. *Nature Neuroscience*, **8**, 1117-1121.
Elman, J. L. (1990). Finding structure in time. *Cognitive Science*, **14**, 179-211.
Elman, J. L., Bates, E. A., Johnson, M. H., Karmiloff-Smith, A., Parisi, D., & Plunkett, K. (1996). *Rethinking Innateness: A connectionist perspective on development*. Cambridge, MA: MIT Press/Bradford Books.
Fodor, J. A., (1983). *Modularity of mind.* Cambridge, MA: MIT Press.
Fodor, J. D., & Frazier, L. (1980). Is the human sentence parsing mechanism an ATN? *Cognition*, **8**, 417-459.
Frazier, L. (1979). *On Comprehending Sentences: Syntactic Parsing Strategies*. Bloomington, Ind.: Indiana University Linguistics Club.
Frazier, L., & Fodor, J. D. (1978). The sausage machine: a new two-stage parsing model. *Cognition*, **6**, 291-325.
Gordon, P.C., Hendrick, R., & Johnson, M. (2001). Memory interference during language processing. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, **27,** 1411-1423.
Hindy, N.C. Altmann, G.T.M., Kalenik, E., and Thompson-Schill, S.L. (under review). Competition between incompatible representations of the same object during event processing.
January D, Trueswell JC, Thompson-Schill SL (2009) Co-localization of Stroop and Syntactic Ambiguity Resolution in Broca's Area: Implications for the Neural Basis of Sentence Processing. *Journal of Cognitive Neuroscience,* **21**, 2434–2444.
Jordan, M. I. (1986). *Serial order: A parallel distributed processing approach* No. 8604). Institute of Cognitive Science, University of California, San Diego.
Kamide, Y., Altmann, G. T. M., & Haywood, S. L. (2003). The time-course of prediction in incremental sentence processing: Evidence from anticipatory eye movements. *Journal of Memory and Language*, **49**, 133–159.
Kimball, J. (1973). Seven principles of surface structure parsing in natural language. *Cognition*, **2**, 15-47.
Kukona, A, Fang, S-Y., Aicher, K.A. Chen, H., & Magnuson, J. (2011). The time course of anticipatory constraint integration. *Cognition*, **119**, 23–42.
Landauer, T. K., & Dumais, S. T. (1997). A solution to Plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, **104**(2), 211-240.
MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994). The lexical nature of syntactic ambiguity resolution. *Psychological Review*, **101**(4), 676-703.
Macleod, C.M. (1991). Half a century of research on the Stroop effect: An integrative review. *Psychological Bulletin*, 109(2), 163-203.

MacWhinney, B. (Ed.). (1999). *The emergence of language*. Mahwah, NJ: Lawrence Erlbaum Associates.
McRae, K., de Sa, V. R., & Seidenberg, M. S. (1997). On the nature and scope of featural representations of word meaning. *Journal of Experimental Psychology: General,* **126**(2), 99-130.
McRae, K., Hare, M.L., Elman, J.E., and Ferretti, T. (2005). A basis for generating expectancies for verbs from nouns. *Memory and Cognition*, **33**, 1174-1184.
Postman, L. (1971). Transfer, interference and forgetting. In J.W. Kling & L.A. Riggs (Eds.), *Woodworth and Schlosberg's Experimental Psychology* (3rd ed), pp. 1019-1132.
Radvansky, G. A., Krawietz, S. A., Tamplin, A. K. (2011). Walking through doorways causes forgetting: Further explorations. *Quarterly Journal of Experimental Psychology*, **8**, 1632-1645.
Rayner, K., Carlson, M., & Frazier, L. (1983). The interaction of syntax and semantics during sentence processing. *Journal of Verbal Learning and Verbal Behaviour**, 22,*** 358-374.
Rogers, T. T., & McClelland, J. L. (2004). *Semantic cognition: A parallel distributed processing approach*. Cambridge, MA: MIT Press.
Steedman, M. (1996). Surface structure and interpretation. Cambridge: MIT Press.
Stroop, John Ridley (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology* **18**(6), 643–662.
Tanenhaus, M. K., Spivey-Knowlton, M. J., Eberhard, K. M., & Sedivy, J. C. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science,* **268**(5217), 1632-1634.
Trueswell, J. C., & Tanenhaus, M. K. (1994). Towards a lexicalist framework of constraint-based syntactic ambiguity resolution. In C. Clifton, L. Frazier, & K. Rayner (Eds.), *Perspectives on Sentence Processing*. (pp. 155-179) Hillsdale, NJ: Lawrence Erlbaum.
Tyler, L. K., & Marslen-Wilson, W. D. (1977). The on-line effects of semantic context on syntactic processing. *Journal of Verbal Learning and Verbal Behaviour*, **16**, 683-692.
Tyler, L. K., & Moss, H. E. (2001): Toward a distributed account of conceptual knowledge*. Trends in Cognitive Sciences*, **5**, 244-252*.*
Valian, V. (1979). The wherefores and therefores of the Competence-Performance distinction. In W. E. Cooper & E. C. T. Walker (Eds.), *Sentence Processing: Psycholinguistic Studies Presented to Merrill Garrett*. (pp. 1-26) New Jersey: Lawrence Erlbaum Associates.
Van Berkum, J.J.A., Brown, C.M., Zwitserlood, P., Kooijman, V., & Hagoort, P. (2005). Anticipating upcoming words in discourse: Evidence from ERPs and reading times. *Journal of Experimental Psychology: Learning, Memory, & Cognition,* **31**, 443-467.
Van Dyke, J. A., & McElree, B. (2006). Retrieval Interference in Sentence Comprehension. *Journal of Memory and Language*, **55**, 157-166.
Wanner, E. (1980). The ATN and the sausage machine: which one is baloney? *Cognition*, **8**, 209-225.
Wanner, E. (1987). The parser's architecture. In F. Kessel (Eds.), *The Development of Language and Language Researchers: Essays in Honour of Roger Brown*. New Jersey: Lawrence Erlbaum Associates.
Wanner, E., & Maratsos, M. (1978). An ATN approach to comprehension. In M. Halle, J. Bresnan, & M. G. A. (Eds.), *Linguistic Theory and Psychological Reality*. (pp. 119-137) Cambridge, MA: MIT Press.
Woods, W. (1970). Transition network grammars for natural language analysis. *Communications of the ACM*., **13**(10), 591-606.
Yee, E., & Sedivy, J. (2006). Eye movements to pictures reveal transient semantic activation during spoken word recognition. *Journal of Experimental Psychology: Learning, Memory and Cognition*, **32**(1), 1-14.
Zaitchik, D. (1990). When representations conflict with reality: the preschooler's problem with false beliefs and false photographs. *Cognition*, **35**, 41–69.