

# LANGUAGE, LINGUISTICS AND COGNITION

Giosuè Baggio, Michiel van Lambalgen, and Peter Hagoort

## 1 INTRODUCTION

Experimental research during the last few decades has provided evidence that language is embedded in a mosaic of cognitive functions. An account of how language interfaces with memory, perception, action and control is no longer beyond the scope of linguistics, and can now be seen as part of an explanation of linguistic structure itself. However, although our view of language has changed, linguistic methodology is lagging behind. This chapter is a sustained argument for a diversification of the kinds of evidence applicable to linguistic questions at different levels of theory, and a defense of the role of linguistics in experimental cognitive science.

### *1.1 Linguistic methodology and cognitive science*

At least two conceptual issues are raised by current interactions between linguistics and cognitive science. One is whether the structures and rules described by linguists are cognitively real. There exist several opinions in this regard, that occupy different positions on the mentalism/anti-mentalism spectrum. At one extreme is cognitive linguistics [Croft and Cruse, 2004], endorsing both theoretical and methodological mentalism. The former is the idea that linguistic structures are related formally and causally to other mental entities. The latter calls for a revision of traditional linguistic methodology, and emphasizes the role of cognitive data in linguistics. At the opposite side of the spectrum lies formal semantics which, partly inspired by Frege's anti-psychologistic stance on meaning and thought [Frege, 1980; Lewis, 1970; Burge, 2005], rejects both versions of mentalism. Somewhere between the two poles is Chomsky's [Chomsky, 1965] theoretical mentalism, which sees linguistic rules as ultimately residing in the brain of speakers. However, his commitment to the cognitive reality of grammar does not imply a revision of linguistic methodology, which is maintained in its traditional form based on native speakers' intuitions and the competence/performance distinction.

The second problem, in part dependent on the first, is whether experimental data on language acquisition, comprehension and production have any bearing on linguistic theory. On this point too, there is no consensus among linguists. The division between competence and performance has often been used to secure linguistics from experimental evidence of various sorts [Bunge, 1984], while

intuitive judgments of native speakers were regarded as the only type of data relevant for the theory [Chomsky, 1965]. However, some authors have granted that at least behavioral data should be allowed to inform competence theories, for instance if the linguist is studying a language which is not her own native language [Marantz, 2005]. Others have proposed frameworks in which competence can be connected to performance mechanisms [Jackendoff, 2002]. But while these models account for how competence constrains performance [Jackendoff, 2007], they seem to overlook the possibility that the reverse is also the case: aspects of linguistic structure may be the outcome of evolutionary processes leading to an adaptation of the brain to language use, that is to performance [Pinker and Jackendoff, 2005]. Generative grammar and formal semantics have regarded accounts of competence as shielded from data provided by experimental psychology and neuroscience. A more inclusive attitude has been adopted by psycholinguists and cognitive brain scientists, driven by an increasing demand of theories and models that would account for their data [Carminati *et al.*, 2000; Featherston *et al.*, 2000; Geurts and van der Slik, 2005; McKinnon and Osterhout, 1996; McMillan *et al.*, 2005]. Despite these attempts, a methodological framework relating linguistics, language processing and low-level neural models is still missing.

## 1.2 *Language, lower and higher cognition*

Most theories in cognitive linguistics and neuroscience regard language as grounded in perception and action. For instance, cognitive semanticists have proposed that the meanings of concrete nouns stored in memory include stereotyped visual-geometric representations of the entities they refer to [Jackendoff, 1987]. Analogously, representations of action verbs might embrace aspects of the relevant motor programs [Hagoort, 1998]. It has also been suggested that the building blocks of semantics like the predicate-argument structure originate in the functional and anatomical organization of the visual and auditory systems [Hurford, 2003]. Experimental work in cognitive neuroscience indicates that language has ties with the sensory-motor systems, but methodology, specific data points and accounts of how exactly language connects to ‘lower’ cognition are still debated [Pulvermuller *et al.*, 2001; Pulvermuller, 2005; Pulvermuller *et al.*, 2005; Ferreira and Patson, 2007; Haslam *et al.*, 2007; Hurford, 2007; Papafragou *et al.*, 2008; Toni *et al.*, 2008; Taylor *et al.*, 2008]. The interested reader may want to follow further these references: in this chapter we will focus on language and ‘higher’ cognitive domains such as planning and reasoning. A motivation for this choice is that planning and reasoning shed light on the computation of complex linguistic structures, which is where language really comes into its own, whereas looking at the interactions between language and sensory-motor systems may give us more insights into representations and processes at the lexical level.

It has been proposed that the recursive organization of plans supplies a mechanism for combinatorial operations in grammar [Steedman, 2002], and the goal-directed nature of planned actions constrains cognitive constructions of time,

causality and events, with consequences for the semantics of tense, aspect and modality [Steedman, 2002; van Lambalgen and Hamm, 2004]. Planning might as well be implicated in the production and comprehension of discourse. Language processing requires adjusting the current discourse model incrementally given the input. If further information counters earlier commitments or expectations, a recomputation of the initial discourse model may be necessary to avoid inconsistencies. This process is best accounted for by the non-monotonic logic underlying planning, and more generally executive function, where the chosen action sequence may be readjusted if obstacles are encountered along the way.

On a par with planning, reasoning is of special interest in this chapter. Some have seen non domain-specific thought and reasoning as the most sophisticated among the cognitive skills subserved by language [Carruthers, 1996; Carruthers, 2002]. This suggestion is sometimes implicit in logical approaches to language since Boole [Boole, 1958, Ch. 2, p. 24] and bears some resemblance to the psycholinguistic notion that reasoning follows and builds upon interpretation [Johnson-Laird, 1980; Singer, 1994]. In this perspective, interpretation equals deriving logical (often classical) form from a sentence's surface structure for subsequent elaboration involving inference. So there is a one-way dependency of reasoning from interpretation: interpretation supports reasoning, though not vice versa. Others have seen the relation between interpretation and inference as based on a two-way interaction [Stenning and van Lambalgen, 2008]: reasoning is involved in computing a model of what is said *and* in deriving conclusions from it. Human communication is thus regarded as the foremost skill enabled by language, and reasoning serves the purpose of coordinating different interpretations of an utterance or different situation models across speakers [Stenning, 2003].

## 2 LINGUISTICS AND COGNITIVE DATA

Let us address in more detail the issues anticipated in section 1.1. In what follows we will present Chomsky's early views on linguistic methodology, introducing a paradox that we believe still lurks in current thinking about relations between linguistics and cognitive data. We will argue that the main problems with the competence/performance distinction are how 'performance' is defined, and what a theory of performance is supposed to include. We will show how this, and the use of intuitions as the only source of evidence in linguistic practice, constitutes an obstacle to a deeper integration of linguistics within cognitive science. These critical sections will be followed by a more constructive part (2.3-2.4), in which Marr's three level scheme is proposed as a replacement of and, we will suggest, an improvement over competence/performance.

### *2.1 A methodological paradox*

It is often said that the relations between cognitive science and linguistics began to be fully appreciated only after the publication of Chomsky's early writings, and

in particular *Aspects of the Theory of Syntax* in 1965. This is certainly true if what is at stake is theoretical mentalism – the notion that linguistic theory deals ultimately with a system of representations and rules in the speaker’s mind/brain. However, although this particular form of theoretical mentalism encourages and to some extent requires some interaction between the two disciplines, the choice of regarding the study of competence as in principle indifferent to the results of experimental research had rather the opposite effect, that of separating theories of meaning and grammar from models of language processing. Many would agree that the contacts between linguistics and cognitive psychology have not been as deep and systematic as they could have been, had various obstacles to fruitful interaction been removed. What is more difficult to appreciate is the existence of a tension in the very foundation of generative grammar, and the inhibiting effect that tension had on the growth of linguistics within cognitive science. Before we move on, it may be worth recovering the terms of this ‘paradox’ directly from Chomsky’s text.<sup>1</sup>

One side of the dilemma is represented by a number of remarks contained in §1 of the first chapter of *Aspects*, where Chomsky writes:

We thus must make a fundamental distinction between *competence* (the speaker-hearer’s knowledge of his language) and *performance* (the actual use of language in concrete situations). Only under [...] idealization [...] is performance a direct reflection of competence. In actual fact, it obviously could not directly reflect competence. A record of natural speech will show numerous false starts, deviations from rules, changes of plan in mid-course, and so on. The problem for the linguist, as well as for the child learning the language, is to determine from the data of performance the underlying system of rules that has been mastered by the speaker-hearer and that he puts to use in actual performance. Hence, in the technical sense, linguistic theory is mentalistic, since it is concerned with discovering a mental reality underlying actual behavior. [Chomsky, 1965, p. 4]

The task of the linguist is that of providing an account of competence based on performance data, that is on normalized records of linguistic behavior. Chomsky grants that performance data are essential to linguistic theorizing. But the issue to be settled, which in fact lies at the heart of the paradox, is exactly what counts as linguistic behavior, or more precisely what kind of performance data can constitute the empirical basis of competence theories. Generative linguists would contend that it was never a tenet of their research program to admit data other than native speakers’ intuitions, but this is not what Chomsky’s remarks suggest. On the contrary, he seems to admit a variety of data types:

---

<sup>1</sup>Over the years Chomsky has entertained different opinions on these issues. Here we choose to focus on those expressed in *Aspects of the Theory of Syntax* [Chomsky, 1965] because these have probably been the most influential. So we identify Chomsky with this particular text rather than with the actual linguist.

Mentalistic linguistics is simply theoretical linguistics that uses performance as data (along with other data, for example, the data provided by introspection) for the determination of competence, the latter being taken as the primary object of its investigations. [Chomsky, 1965, p. 193]

The evidential base of linguistics consists of introspective judgments *and* performance data, that Chomsky mentions here as if they were in an important sense different from intuitions. Moreover, intuitions are alluded to here as a subsidiary source of evidence, and as part of a larger class of data types. The question is precisely what should be considered performance data. Is elicited and experimentally controlled behavior allowed to exert some influence on accounts of competence? There are reasons to believe that Chomsky would have answered affirmatively, the most important of which has to do with his remarks on the limits of intuitions. In 1955, in *The Logical Structure of Linguistic Theory* [Chomsky, 1955] he wrote:

If one of the basic undefined terms of linguistic theory is ‘intuition’, and if we define phonemes in this theory as elements which our intuition perceives in a language, then the notion of phoneme is as clear and precise as is ‘intuition’. [...] It should be clear, then, why the linguist interested in constructing a general theory of linguistic structure, in justifying given grammars or (to put the matter in its more usual form) in constructing procedures of analysis should try to avoid such notions as ‘intuition’. [Chomsky, 1955, pp. 86-87]

An even more explicit position was expressed in the 1957 book *Syntactic Structures*, where Chomsky suggests that hypotheses on properties of linguistic strings and their constituents should be evaluated on the basis of controlled operational tests. Relying on native speaker’s judgments or intuitions, he writes,

amounts to asking the informant to do the linguist’s work; it replaces an operational test of behavior (such as the pair test) by an informant’s judgment about his behavior. The operational tests for linguistic notions may require the informant to respond, but not to express his opinion about his behavior, his judgment about synonymy, about phonemic distinctness, etc. The informant’s opinions may be based on all sorts of irrelevant factors. This is an important distinction that must be carefully observed if the operational basis for grammar is not to be trivialized. [Chomsky, 1957, pp. 8-9]<sup>2</sup>

Controlled operational tests are thus necessary in order to overcome the difficulties arising from relying exclusively on native speakers’ intuitions. This implies that

<sup>2</sup>The circularity which Chomsky is alluding to here is also mentioned by Quine in his 1970 paper on linguistic methodology: “We are looking for a criterion of what to count as the real or proper grammar, as over against an extensionally equivalent counterfeit. [...] And now the test suggested is that we ask the native the very question we do not understand ourselves: the very question for which we ourselves are seeking a test. We are moving in an oddly warped circle.” [Quine, 1970, p. 392].

introspective data are dismissed as an inadequate source of evidence for linguistic theory. So here is one horn of the dilemma: mentalistic linguistics rejects speakers' intuitions and requires performance data, including controlled behavioral tests, to constrain the theory of competence.

The other side of the paradox is represented by a series of remarks in §4 of chapter 1 of *Aspects*, where Chomsky questions the nature of the empirical basis of competence theories:

There is, first of all, the question of how one is to obtain information about the speaker-hearer's competence, about his knowledge of the language. Like most facts of interest and importance, this is neither presented for direct observation nor extractable from data by inductive procedures of any known sort. Clearly, the actual data of linguistic performance will provide much evidence for determining the correctness of hypotheses about underlying linguistic structure, along with introspective reports (by the native speaker, or the linguist who has learned the language). [Chomsky, 1965, p. 18]

Experimental research based on controlled observation and statistical inference is seen as providing facts of no 'interest and importance', and rejected as ineffective for the purposes of the theory of competence. Interestingly, intuitions are treated as if they were on a par with performance data. Not for long, however, because Chomsky a few paragraphs later takes an important step away from psychology:

The critical problem for grammatical theory today is not a paucity of evidence but rather the inadequacy of present theories of language to account for masses of evidence that are hardly open to serious question. The problem for the grammarian is to construct a description and, where possible, an explanation for the enormous mass of unquestionable data concerning the linguistic intuition of the native speaker (often, himself); the problem for one concerned with operational procedures is to develop tests that give the correct results and make the relevant distinctions. [...] We may hope that these efforts will converge, but they must obviously converge on the tacit knowledge of the native speaker if they are to be of any significance. [Chomsky, 1965, pp. 19-20]

The range of data that could affect the theory of competence has been narrowed down to intuitions, and more specifically to those of the linguist. The task of experimental research, Chomsky says, is to develop tests that would ultimately align with introspective data. The convergence of linguistics and psychology is thus projected forward in time as a desirable outcome not of the joining of efforts, but of their strict segregation. Not only are linguistics and psychology now regarded as separate enterprises, but psychology is also required – in order to meet a standard of explanatory adequacy – to provide results that are consistent with the theory of competence as based on the linguist's intuitions. The second horn of the dilemma

is thus the following: linguistic theory is based primarily on the intuitions of native speakers, and does not require controlled experimentation to constrain accounts of competence.

## 2.2 *The vagaries of intuition*

For some linguists, in particular in generative grammar and formal semantics, the intuitions of native speakers constitute the empirical basis of the theory of competence. But the prominent place assigned to intuitions by modern linguistic methodology seems at odds with the principles of mentalism. If competence is a system of rules and structures realized in the speaker's brain, and if behavior reflects the functioning of such system, then a linguist constructing a competence theory – and perhaps analogously a child learning a language – must solve an 'inverse problem', that of inferring the rules of competence from observable performance. In order to solve this problem, the linguist might need to take into account a broad range of data. So *any* reliable physiological or behavioral measure of performance should, at least in principle, be allowed to contribute to the theory of competence. The question is where should one draw a line between relevant (intuitions?) and irrelevant (neurophysiology?) data, and why. Until convincing answers are found, it would seem that the more comprehensive one's methodological framework, the better. Here is why mentalism is to be preferred over traditional philosophies of language.

The conflict with mentalism is however not the only problem raised by introspective judgments. Another source of concern is Chomsky's claim that intuitions are not only the starting point of linguistic theorizing, but also the standard to which any grammar should conform:

A grammar can be regarded as a theory of language; it is *descriptively adequate* to the extent that it correctly describes the intrinsic competence of the idealized native speaker. The structural descriptions assigned to sentences by the grammar, the distinctions that it makes between well-formed and deviant, and so on, must, for descriptive adequacy, correspond to the linguistic intuition of the native speaker (whether or not he may be immediately aware of this) in a substantial and significant class of crucial cases. [Chomsky, 1965, p. 24]

Supposing the tension with mentalism were relieved, allowing other data types to influence competence models, and introspective judgments were used only at the outset of linguistic inquiry, intuitions would still pose a number of serious methodological problems. It is not just the role of intuitions in linguistic theorizing that has to be put under scrutiny, but also the claim that intuitions offer a vantage point on tacit knowledge.

### 2.2.1 *Intuitions in linguistics*

If the system of linguistic rules in a speaker's brain really is "deeply unconscious and largely unavailable to introspection" [Jackendoff, 2003, p. 652], one should see discrepancies between overt linguistic behavior, that reflects 'unconscious' competence rules, and the intuitions or beliefs that speakers have about these rules. This notion has been substantiated by Labov [Labov, 1996], who collected evidence on a wide variety of cases in regional American English. One example is the positive use of 'anymore' in various sections of the Philadelphia white community, meaning that a situation which was not true some time in the past is now true, roughly equivalent to 'nowadays':

- (1) Do you know what's a lousy show anymore? Johnny Carson.

Labov interviewed twelve speakers who used the adverb freely and consistently with its vernacular meaning exemplified in (1). He reported a majority of negative responses when they were asked whether a sentence like (1) is acceptable, and surprisingly weak intuitions on what the expression signifies in their own dialect, which contexts are appropriate for its use, and what inferences can be drawn from its occurrences.

Other arguments suggest that the use of intuitions in linguistics is problematic in many ways. For instance, [Marantz, 2005] has observed that grammaticality is a technical term defined within linguistic theories: a sound/meaning pair is grammatical or well-formed with respect to a grammar if and only if that grammar generates or assigns a structural description to the pair such that all relevant grammaticality or well-formedness constraints can be satisfied. In the quote from *Aspects* above, Chomsky takes for granted that structural descriptions assigned by some grammar to sentences can be checked for correspondence against native speakers' judgments. However, native speakers of a language can hardly be said to have intuitions of grammaticality in the technical sense, nor can they grasp other properties of strings as these are defined within a formal grammar. Moreover, naïve language users might conflate into the notion of grammaticality different morpho-syntactic, semantic and pragmatic criteria of acceptability, and they might do so in a way that is beyond control for the linguist. Similar observations would also apply to intuitive judgments of synonymy or truth-conditions, as opposed to formal definitions within a semantic theory.

As a way out, one might argue that a caveat only applies to naïve informants, and that the intuitions of linguists, immune to pre-theoretical notions of grammaticality, synonymy, and the like, are in fact reliable [Devitt, 2006]. Relevant to this issue, is an experiment by [Levelt, 1972] in which the intuitions of twenty-four trained linguists were investigated. Participants were presented with fourteen examples from their own field's literature, among which:

- (2) a. No American, who was wise, remained in the country.  
 b. The giving of the lecture by the man who arrived yesterday assisted us.



None of the linguists rated correctly the ungrammatical sentence (2a), and sixteen judged the well-formed sentence (2b) as ungrammatical. Ungrammatical sentences had less chance of being judged ungrammatical than grammatical items. Levelt warns against taking these results too seriously, but he observes with some reason that “they are sufficiently disturbing to caution against present day uses of intuition” [Levelt, 1972, p. 25].

We could go on providing other examples of the problems that might arise with the use of introspective reports in the analysis of specific natural language sentences. However, we should now like to take a different approach, considering an argument targeted at the nature and scope of intuitions. The argument, introduced and discussed by Hintikka [Hintikka, 1999], starts with the observation that intuitions of grammaticality, synonymy etc. always relate to particular sentences (i.e. tokens), and not to entire classes of items, or to the common syntactic or semantic structure they share (i.e. types). Hintikka writes that

intuition, like sense perception, always deals with particular cases, however representative. [...] But if so, intuition alone cannot yield the general truths: for instance, general theories for which a scientist and a philosopher is presumably searching. Some kind of generalizing process will be needed, be it inductive inference, abduction, or a lucky guess. The intuitions [Chomsky] recommended linguists to start from were intuitions concerning the grammaticality of particular strings of symbols, not concerning general rules of grammar. [Hintikka, 1999, p. 137-138]

Against Hintikka’s claim, one may argue that also paradigmatic variation is a proper object of intuition. The linguist would then be able to generalize over the properties of linguistic structures by constructing a paradigmatic set of sentences exhibiting those properties. This view however can be countered with the observation that the supposed ‘intuitions’ about paradigmatic cases are more similar to theory-laden hypotheses than to introspective judgments of naïve informants. The linguist, in order to construct such paradigmatic items, has to be able to control all irrelevant variables and systematically manipulate the factors of interest. This, in turn, requires that the linguist knows details of the grammar or the logical structure of the language which seem inaccessible to naïve speakers. It is this knowledge, which is often drawn from existing theories, that allows the linguist to have intuitions about linguistic structure. This leads us to Hintikka’s key statement, that competence theories are not equipped with built-in devices for deriving abstract grammatical or semantic forms from particular linguistic samples. That is,

reliance on generalization from particular cases is foreign to the methodological spirit of modern science, which originated by looking for dependencies of different factors in instructive particular cases (often in controlled experimental situations), and by studying these dependences by the same mathematical means as a mathematician uses in studying

the interdependencies of different ingredients of geometrical figures in analytic geometry. [...] transformational grammarians and other contemporary linguists would do a much better job if, instead of relying on our intuitions about isolated examples, they tried to vary systematically suitable ingredients in some sample sentence and observed how our ‘intuitions’ change as a consequence. Now we can see why such systematic variation is a way of persuading our intuitions to yield general truths (dependence relations) rather than particular cases. [Hintikka, 1999, p. 135]

If intuitions are to serve as a reliable starting point in linguistic inquiry, they should be proved to have systematic properties. Observing patterns of covariation of introspective judgments and other factors – such as the structure, the content, the context of occurrence of the sentence, the attitude of the speaker, and so on – would make the particular examples under consideration instructive and thus effective as part of the empirical basis of linguistic theories. The important consequence is that, in order to systematically alter the ingredients of sample sentences, the linguist should be able to control these factors in a manner similar to the manipulation of experimental variables in laboratory research. The solution offered by Hintikka to the problem of intuitions points in the direction of *infusing linguistic practice with psychological experimentation*. The linguist would as usual start from intuitions, but only the systematic aspects of these as revealed by experimentation, and if necessary statistical tests, would be preserved and transferred into the theory (see [Bunge, 1984, pp. 158-163] for a similar position).<sup>3</sup> Hintikka offers an intriguing example, in which one tries to define the meaning of an expression in Montague grammar on the basis of systematic dependencies between subjects’ intuitions and the contexts of occurrence of the expression of interest. In particular, he writes, if the notion of possible world is allowed in the theory,

then there is, in principle, no definite limit as to how far your experimentation (construction of ever new situations) can carry you in determining the class of scenarios in which the word does or does not apply. And such a determination will, at least for a Montagovian semanticist, determine the meaning of the word. Indeed, in Montague semantics, the meaning of a term is the function that maps possible worlds on references (extensions) of the appropriate logical type (category). And such functions can, in principle, be identified even more and more fully by systematic experimentation with the references that a person assigns to his terms in different actual or imagined scenarios.

---

<sup>3</sup>[Bunge, 1984, p. 168] pinpoints several methodological choices in generative linguistics which seem to diminish its relevance in empirical science, such as the “conduct of linguistic inquiry in total independence from neuroscience, social science, and even scientific psychology” and “a heavy reliance on intuition”. We too consider these as obstacles to understanding language, but we disagree with the judgment that Bunge formulates based on these remarks – that modern linguistics is (or has been) pseudo-scientific.

[Hintikka, 1999, p. 146]<sup>4</sup>

However, it may be a fair point in favor of introspective judgments in a broader sense to add that Hintikka considers thought experiments on a par with genuine experimentation [Hintikka, 1999, p. 146]. Thus, instead of eliciting overt responses from subjects in a number of conditions, the experimenter imagines herself in such situations. If the relevant variables are controlled with as much care as one would exercise in an experimental setting, introspection can reveal systematic aspects of language use, and thus contribute to theories of competence.

Hintikka's argument can be made more explicit with reference to a number of studies investigating the role of the most important of his 'suitable ingredients' – context. Linguistic and psycholinguistic research has demonstrated that the context in which a sentence occurs can affect judgments of acceptability. [Bolinger, 1968] reported that sentences, which speakers judge as semantically implausible when presented in isolation, are regarded as acceptable when embedded in context. Consider the following examples:

- (3) a. It wasn't dark enough to see.  
b. I'm the soup.

These sentences are typically judged as semantically deviant, although for different reasons: (3a) because one normally needs light in order to see, and (3b) because the predicate 'being a soup' cannot be applied to a human being. Now, consider the same sentences embedded in a suitable discourse context, with (4b) being spoken at a cashier's counter in a restaurant:

- (4) a. I couldn't tell whether Venus was above the horizon. It wasn't dark enough to see.  
b. You've got us confused. You're charging me for the noon special. The man in front of me was the noon special. I'm the soup.

Examples (3) in an appropriate context seem perfectly acceptable. Because context has such marked effects on intuitions, linguistic theory, if it has to rely on introspective judgments, should explicitly take into account this fact.

### 2.2.2 *Intuitions in psycholinguistics*

The appeal to intuitions was not an explicit choice of methodology in psycholinguistics and the cognitive neuroscience of language. In fact, the method of introspection was discarded in scientific psychology after its failures in the 19th century. However, it is adopted in language processing research as a means of establishing differences between sentence types to be used as stimuli in actual experiments.

---

<sup>4</sup>Although we consider Hintikka's an informative example of linguistic theorizing based on covariation patterns of contextual factors and intuitions, we must also add that there are serious problems with the notion of meaning (that is, Frege's *Sinn*) in Montague semantics. For instance, since the theory allows for infinitely many possible worlds, it becomes unclear whether we can even approximate the meaning of an expression using Hintikka's method.

The typical procedure for setting up a language processing study is to start with a few sample sentences differing with respect to some linguistic feature, the assessment of which is initially left to the intuitions of the experimenter. For instance, let us consider one of the first ERP studies on syntactic constraint violations, by [Osterhout and Holcomb, 1992]. Here the starting point is a pair – or a relatively small set of pairs – of sentences containing either an intransitive (5a) or a transitive (5b) verb, using a direct object construction:

- (5) a. The woman struggled to prepare the meal.  
 b. The woman persuaded to answer the door.

Up to this stage, the methodology is by and large the same as that of the linguist. However, while the latter would then proceed with, say, formalizing the requirements of intransitive and transitive verbs with respect to direct objects, the psycholinguist, to make sure there is sufficient statistical power to test her processing hypotheses in a dedicated experiment, would have to construct a large set of sentences with the same structure and properties of (5a-b). In the next step, the sentences would be presented to subjects while the dependent variables of interest are measured, which in the study of Osterhout and Holcomb were ERPs and grammaticality ratings. Grammatical sentences like (5a) were judged to be acceptable in 95% of the cases, and supposedly ungrammatical items like (5b) in 9% of the cases. One may argue, as an academic exercise towards an explanation of the 9% figure, that (5b) does have contexts in which it is both grammatical and semantically acceptable, for instance if it is interpreted as a reduced relative clause ('The woman who was persuaded to answer the door'), and is uttered as an answer to a *who* question, as in the following dialogue:

- (6) A: Who stumbled on the carpet in the hallway?  
 B: The woman persuaded to answer the door.

We have already encountered this phenomenon discussing Bolinger's examples above. In a context such as (6), Osterhout and Holcomb's ungrammatical sentence becomes perfectly admissible. Acceptability judgments, therefore, depend on the range of uses (or contexts of use) readers are willing to consider. In this sense, subjects' intuitions may differ from those of the experimenter. For example, a linguist would remind us that 'The woman persuaded to answer the door' is an NP, and not a sentence. But what prevents naïve language users from including well-formed NPs into their notion of 'sentence'? Here the answer can only be: the linguist's own notion of 'sentence'. This also suggests that discrepancies between the intuitions of naïve informants and trained scientists may be more important than isolated linguists' intuitions when it comes to fully explaining a data set.

### 2.3 *Beyond competence and performance*

Intuitions are but one of the many sources of concern for a thoroughly mentalistic approach to language. As [Jackendoff, 2002, p. 29] has pointed out, there is

a conflict, which roughly coincides with the dilemma as we described it above, between Chomsky's theoretical mentalism and traditional linguistic methodology as based on intuitions and on the competence/performance distinction. Mentalism requires at least a revision of that distinction. [Jackendoff, 2002] has addressed this issue, trying to find a more accommodating formulation which allows a natural interplay of linguistics and the cognitive sciences:

Chomsky views competence as an idealization abstracted away from the full range of linguistic behavior. As such, it deserves as much consideration as any idealization in science: if it yields interesting generalizations it is worthwhile. Still, one can make a distinction between 'soft' and 'hard' idealizations. A 'soft' idealization is acknowledged to be a matter of convenience, and one hopes eventually to find a natural way to re-integrate excluded factors. A standard example is the fiction of a frictionless plane in physics, which yields important generalizations about forces and energy. But one aspires eventually to go beyond the idealization and integrate friction into the picture. By contrast, a 'hard' idealization denies the need to go beyond itself; in the end it cuts itself off from the possibility of integration into a larger context.

It is my unfortunate impression that, over the years, Chomsky's articulation of the competence/performance distinction has moved from relatively soft [...] to considerably harder. [Jackendoff, 2002, p. 33]

Jackendoff suggests we adopt a 'soft' competence/performance distinction, adding a third component to the framework [Jackendoff, 2002]. The theory of competence is seen as the characterization of phonologic, syntactic and semantic data structures as they are processed and stored in the brain of speakers during language acquisition. The theory of performance is the description of the actual occurrence of such data structures in language comprehension and production. The theory of neural instantiation is an account in terms of brain structures and processes of competence and performance. Jackendoff provides an architecture in which competence components (phonology, syntax and semantics, plus interface rules) interact in a manner that is consistent with the incrementality and the 'opportunism' (his label for immediacy) of language processing [Jackendoff, 2007]. However, to solve the dilemma described above, it is not enough to show that competence determines the state-space available to users of a language during performance [Jackendoff, 2002, p. 56]. The issue is, rather, whether there is interplay between competence and performance, that is – turning Jackendoff's tag line upside down – whether the logic of processing dictates the logic of competence, and to what extent.

As we saw above, in his early writings Chomsky claimed that theories of competence have nothing to learn from processing data [Chomsky, 1965]. Minimalists have suggested that syntax is adapted to the requirements holding at the interface with other cognitive modules, such as the sensory-motor and conceptual systems. However, they deny what functionalists on the contrary accept, namely

that syntax is well-designed for use [Chomsky *et al.*, 2002; Hauser *et al.*, 2002; Fitch *et al.*, 2005]. Evidence against adaptation to performance is provided, according to minimalists, by memory limitations, constructions such as garden-path and center embedding sentences, which seem suboptimal in various respects. Here two remarks are in order. The first is that such phenomena do not constitute evidence against adaptation *per se*, but rather (if anything like that is possible) against ‘perfect’ adaptation. Minimalists seem to commit what optimality theorists have called the ‘fallacy of perfection’ [McCarthy and Prince, 1994], consisting in confusing optimal outcomes, which are the result of equilibria between different variables, with best possible outcomes for just a subset of the factors involved, for instance the absence of unstable or ambiguous constructions (see [Pinker and Jackendoff, 2005] for discussion). The second remark is that, even if we assume that competence is neither perfectly nor optimally adapted to use, it still seems conceivable that performance constraints shaped competence rules. Therefore, the problem is not whether language is an adaptation: that some traits of competence reflect the outcomes of adaptive evolutionary processes acting on actual brain systems, including adaptation to communication needs, seems to be a widely accepted view [Hauser *et al.*, 2002; Pinker and Jackendoff, 2005; Fitch *et al.*, 2005]. The problem is rather: (how) can we construct a methodological framework in which it is possible to determine what aspects of competence can be explained adaptively?

The reason why generative linguistics does not seem capable of addressing this issue is, in our opinion, to be attributed more to how performance is defined than to a rigid view of the competence/performance distinction. [Jackendoff, 2002, p. 30] rightly observes that in Chomsky’s original proposal a large and heterogeneous set of phenomena were collapsed into ‘performance’: errors, shifts of attention, memory limitations, processing mechanisms, and so on. Only a very superficial assessment of the factors involved in language use could justify the notion that a single, relatively compact theory of performance could account for all those phenomena. It seems more reasonable to assume that different theories, developed using different analytical approaches, are necessary to understand how language interacts with memory and attention, how errors of different type and origin are produced (for also language disorders give rise to performance failures), and so on. We agree with Jackendoff on the characterization of competence and neural implementation, but we believe a more appropriate intermediate level should be chosen.

#### 2.4 *Marr’s three-level scheme as applied to language*

Jackendoff’s updated view of the competence/performance distinction as a soft methodological separation, plus a theory of neural realization, resembles Marr’s 1982 tripartite scheme for the analysis of cognitive systems [Spivey and Gonzalez-Marquez, 2003]. Marr suggested that cognitive processes should be modeled at three, nearly independent levels of analysis: a computational level (what is com-

puted?), an algorithmic level (how is computation carried out?), and a level of physical implementation (what are the properties of the real machines that can execute the algorithms?). From this perspective, Jackendoff's trajectory away from Chomsky appears incomplete. There is a partial redefinition of competence/performance, and the addition of a third level, but it is doubtful whether this move leads to the kind of transitions and mutual constraints between levels of analysis afforded by Marr's scheme. So it may be worth asking what would be the advantages of taking one step further, that is replacing the competence/performance distinction with Marr's distinction between computational and algorithmic analyses.

An important consequence of this choice is that performance theory is now seen as an intermediate level of analysis at which the algorithms and memory mechanisms supporting specific linguistic computations are described. That might seem a rather abrupt move, as it restricts the scope of performance to algorithms, and thereby leaves aside a large number of phenomena which, some might suggest, cannot be adequately treated in algorithmic terms. For instance, conscious inner speech is an important function of language [Carruthers, 1996; Carruthers, 2002], and one in which there seems to be no definite input-output mapping involved. On the other hand, for those phenomena that are best treated as structured input-output processes, for example language production and comprehension, Marr's framework allows competence theories, if properly construed, to be investigated as part of actual information processing systems. Applications of this idea to semantics will be shown below.

Does our appeal to Marr's three-level scheme solve the problems associated with the competence/performance distinction? It seems it does, because the variety of phenomena that were collapsed into performance can now be understood in their distinctive features. For instance, working memory as involved in a given task can be examined at the level of algorithms. The algorithmic analysis may suggest a description of the memory architecture and the memory resources required by the system, and this constitutes a first step toward an explanation in neural terms. Conversely, memory organization constrains the classes of algorithms that can be computed by that machine, and the type of data structures that the computational theory can produce. An example of bottom-up adjustment is Yngve's 1960 explanation of the dominance of right-branching over left-branching and center-embedding structures. Another example are 'minimal models' of discourse, as we shall see later.

In brief, one key feature of Marr's scheme is that the way levels of analysis are defined makes it easier to see how formal theories of grammar, language processing and neural computation could be integrated and mutually constrained. It seems that, preserving the notion of 'performance', and a fortiori 'competence', such well-constrained formal routes between levels of analysis would become less accessible.

Below we apply this tentative methodological sketch to the semantics of tense, aspect and event structure. Our goal is to devise semantic analyses that are formally specified and cognitively motivated, that is, highlighting connections be-

tween the meanings of temporal expressions, planning and reasoning. The semantic analyses should also be algorithmically explicit, such that processing predictions, or general constraints on processing architecture, can be formulated. We hope to show that our theory of tense, aspect, and event structure not only meets these requirements, but can also be used to provide alternative explanations of existing experimental data on language comprehension. The last part of this chapter puts our enterprise into a wider neuroscience-oriented perspective, introducing the ‘binding problem for semantics’.

### 3 PLANNING, REASONING, MEANING

#### 3.1 *The cognitive substrates of tense and aspect*

We see it as the essential purpose of tense and aspect to facilitate the computation of event structure as described in a narrative. One consequence of this characterization is that, contrary to what generative and formal semantic approaches maintain, it is not very useful to study tense and aspect at the sentence level only. Tense, aspect and event structure really come into their own only at the discourse level [Comrie, 1976; Comrie, 1985]. Tense and aspect, however, cannot by themselves determine event structure, and must recruit world knowledge. Examples (7a-c) will make clear what we have in mind.

French has several past tenses (Passé Simple, Imparfait, Passé Composé), which differ in their aspectual contributions. The following mini-discourses in French<sup>5</sup> all consist of one sentence in the Imparfait and one in the Passé Simple. However, the structure of the set of events differs in each case.

- (7) a. Il faisait chaud. Jean ôta sa veste. (Imp, PS)  
*It was hot. Jean took off his sweater.*  
 b. Jean attrapa une contravention. Il roulait trop vite. (PS, Imp)  
*Jean got a ticket. He was driving too fast.*  
 c. Jean appuya sur l'interrupteur. La lumière l'éblouissait. (PS, Imp)  
*Jean pushed the button. The light blinded him.*

In the first example, the Imp-sentence describes the background against which the event described by the PS-sentence occurs. In the second example, the PS-event terminates the Imp-event, whereas in the third one the relation is rather one of initiation. These discourses indicate that world knowledge in the form of knowledge of causal relations is an essential ingredient in determining event structure. This knowledge is mostly applied automatically, but may also be consciously recruited if the automatic processing leaves the event structure underdetermined. It is the task of cognitive science to determine what this algorithm looks like, and how

---

<sup>5</sup>Taken from an unpublished manuscript on French tenses by [Kamp and Rohrer, 1985]. See also [Eberle and Kasper, 1989] and [Asher and Lascarides, 2003].



it is actually implemented. The prominent role of causal relationships in (7a-c) suggests a direction in which to look for that algorithm.<sup>6</sup>

### 3.2 *Planning, causality and the ordering of events*

Stated bluntly, our hypothesis is:

The ability to automatically derive the discourse model determined by a narrative is subserved by the ability to compute plans to achieve a goal.

At first this hypothesis may seem unintelligible: what do goals and plans have to do with discourse? But as we will see, it is possible, indeed advantageous, to represent tense and aspect formally as goals to be satisfied. A discourse then sets up a system of interlocking goals, which is at least formally similar to the complex goals that occur in, say, motor planning. The hypothesis then says that the formal similarity arises from the fact that the very same cognitive mechanism is responsible for dealing with goals in both domains, motor control and language processing.

Here we present several converging lines of evidence which lend some plausibility to this conjecture. Firstly, at a very general level one can say that a distinguishing feature of human cognition is that it is goal-oriented, with goals that range from very short-term (get a glass of water) to very long-term (having sufficient income after retirement). In each case, the goal is accompanied by a plan which produces an ordered collection of actions, be they motor actions or transfers of money to a special account. More precisely, planning consists in

the construction of a *sequence*<sup>7</sup> of actions which will achieve a given goal, taking into account properties of the world and the agent, and also events that might occur in the world.

Given the ubiquity of goal-plan organisation in human cognition, it is not surprising that there have been numerous attempts to link the language capacity with the planning capacity. The setting is usually a discussion of the origins of language. Even if it is granted that some non-human primates have learned a primitive form of language, there is still a striking difference in language proficiency between chimpanzees and ourselves. It is still a matter of ongoing debate to determine exactly what this difference consists in. Some would say that the difference is in the syntax: human syntax is recursive, the chimpanzee's syntax (if that is the word) is not. One may then point to an analogy between language and

---

<sup>6</sup>There is a body of literature on what are called 'situation models' (what we have called 'discourse models' or 'event structures') which contains suggestive evidence to show that these models not only represent objects and events introduced by the discourse, but also general and specific causal information about the world not explicitly mentioned in the discourse. Space constraints forbid extensive discussion of this line of research; we can only direct the reader to the survey article [Zwaan and Radvansky, 1998].

<sup>7</sup>More complex plans are possible, involving overlapping actions.

planning. Language production can be characterized as transforming a semantic structure, to which the notion of linearity may not be applicable, into linear form, that is an utterance. Planning also involves linearization, and that is how the language-planning connection is drawn. An alternative strategy, not inconsistent with the first, is to show that the recursive structure of syntax is linked to the recursive structure (or hierarchical organization) of plans [Greenfield, 1991; Steedman, 2002]. Non-human primates engage in planning for time spans not exceeding 48 hours, as is known since Kohler's 1925 observations. This has also been attested in squirrel monkeys in experiments by [McGonigle *et al.*, 2003]. What these experiments jointly show is that on the one hand planning provides a link between humans and non-human primates, but that on the other hand complex planning sets humans apart from non-human primates. As such the planning capacity can be a starting point for discussions of the origins of language because it may account for both continuity (capacity for planning shared between humans and their ancestors) and change (increased capacity for planning leading to human linguistic abilities).

A more direct connection between language and planning, and one focussing on semantics rather than syntax, was investigated experimentally by [Trabasso and Stein, 1994] in a paper whose title sums up their program: *Using goal-plan knowledge to merge the past with the present and the future in narrating events on-line*. Trabasso and Stein argue that “the plan unites the past (a desired state) with the present (an attempt) and the future (the attainment of that state)” [Trabasso and Stein, 1994, p. 322], “[c]ausality and planning provide the medium through which the past is glued to the present and future” [Trabasso and Stein, 1994, p. 347]. They present the results of a study in which children and adults were asked to narrate a sequence of 24 scenes in a picture storybook called *Frog, where are you?*, in which a boy tries to find his pet frog which has escaped from its jar.<sup>8</sup> The drawings depict various failed attempts, until the boy finds his frog by accident. The aim of the study is to determine what linguistic devices, in particular temporal expressions, children use to narrate the story as a function of age. The authors provide some protocols which show a child of age 3 narrating the story in a tenseless fashion, describing a sequence of objects and actions without relating them to other objects and actions. None of the encoded actions is relevant to the boy's ultimate goal. Temporal sequencing comes at age 4, and now some of the encoded actions are relevant to the goal. Explicit awareness that a particular action is instrumental towards the goal shows up at age 5. At age 9, action-goal relationships are marked increasingly, and (normal) adults structure the narrative completely as a series of failed or successful attempts to reach the goal. Thus we see that part of what is involved in language learning is acquiring the ability to produce discourse in such a way that a goal-plan structure is induced in the hearer.

---

<sup>8</sup>This is a classic experimental paradigm for investigating the acquisition of temporal notions in children. See Berman and Slobin [Berman and Slobin, 1994] for methods, results and, last but not least, the frog pictures themselves. We will come back to this paradigm when discussing use of verb tense in children with ADHD in section 3.2.3.

The authors' claim is that such discourse models are never mere enumerations of events, but that our very mental representation of time privileges discourse models in which events can be viewed as part of a plan.

Our proposal is that also when viewed computationally, discourse models are best treated as plans, i.e. as output of the planning mechanism. Indeed, it is of some interest to observe that the ingredients that jointly enable planning have a prominent role to play in the construction of a discourse model. Take for instance causality, shown to be involved in the interpretation of (7a-c). Planning essentially requires knowledge of the causal effects of actions as well as of the causal effects of possible events in the world. Accordingly, the planning capacity must have devised ways of retrieving such knowledge from memory. Planning also essentially involves ordering actions with respect to each other and to events occurring in the world which are not dependent upon the agent. Furthermore, the resulting structure must be held in memory at least until the desired goal is attained. The reader can easily envisage this by considering the planning steps that lead to a pile of pancakes. For instance, causal knowledge dictates that one has to pour oil in the frying-pan before putting in the batter, and this knowledge has to remain active as long as one is not finished.

While the preceding considerations point to some data structures common to both plans and discourse models, the fundamental *logical* connection between discourse processing and planning is that both are non-monotonic. When we plan, deliberately or automatically, we do so in virtue of our best guess about the world in which we have to execute our plan. We may plan for what to do if we miss the bus, but we don't plan for what to do if the bus doesn't come because the gravitational constant changes, even though that is a logical possibility. Similarly, the computation of a discourse structure may be non-monotonic. For instance, the reader who sees (8a) is likely to infer (that is, to read off from the discourse model) that Bill is no longer a member, but that implicature can easily be canceled, as in (8b):

- (8) a. Bill used to be a member of a subversive organization.  
 b. Bill used to be a member of a subversive organization, and he still is.

The discourse model belonging to (8b) is not simply an extension of the one for (8a), although (8b) is an extension of (8a); but the temporal interpretation of the main clause must be recomputed in going from (8a) to (8b). We will see more examples of this phenomenon when investigating the relation between verb tenses and planning.

We propose that the link between planning and temporal semantics is provided by the notion of *goal*. In both comprehension and production, the goal interpreting the tensed VP is to introduce the event corresponding to the tensed VP into the already existing event structure. This goal always has two components:

1. location of event in time;
2. meshing it with other events.

The role of planning is to establish definite temporal relationships between the events involved.

How planning can do this is best illustrated by means of an example. Consider what goes on in comprehending

(9) Max fell. John pushed him.

On one prominent reading, the event described in the second sentence precedes, indeed causes, that described in the first sentence. The relevant goals are in this case:

update discourse with past event  $e_1 = fall(m)$  and fit  $e_1$  in context  
 update discourse with past event  $e_2 = push(j, m)$  and fit  $e_2$  in  
 context

Planning must determine the order of  $e_1$  and  $e_2$ , and to do so the planning system recruits causal knowledge as well as the principle that causes precede effects. We sketch an informal argument here; the next subsection gives more formal details, necessary to understand the connection between discourse processing viewed as non-monotonic computation, and traces of this computation in neural responses. The informal argument runs like this. There is (assumed to be) no context yet in which  $e_1$  must be processed, so  $e_1$  is simply located in the past. When it comes to interpreting  $e_2$ , we have a context ( $e_1$ ). The planning mechanism now retrieves possible relationships involving both  $e_1$  and  $e_2$ , and one of these is that a push initiates falling. Since the cause comes before its effect, this yields that  $e_2$  precedes  $e_1$ . Observe that this is a default inference only; as we will see, it is possible to unify  $e_1$  and  $e_2$  with other material such that their temporal order becomes reversed.

### 3.2.1 Computations on event structures<sup>9</sup>

To give the reader a detailed picture of what goes on in such computations, we have to introduce some notation, borrowed from the Event Calculus [van Lambalgen and Hamm, 2004], which will also be useful for our discussion of the ‘binding problem’ later in this chapter. We make a distinction between *events* (denoted  $e, e', \dots, e_0, \dots$ ) and *processes* or *fluents* (denoted  $f, f', \dots, f_0, \dots$ ). We say that events occur or *happen* at a particular time, and represent this by the expression  $Happens(e, t)$ . By contrast, processes do not occur, but are going on at a particular time, and for this we use the predicate  $HoldsAt(f, t)$ . Events and processes can stand in causal relations. For instance, an event may kick off a process:  $Initiates(e, f, t)$ ; or it may end one:  $Terminates(e, f, t)$ . We will use these predicates to mean the causal relation only. It is not implied that  $e$  actually occurs. Finally, a useful predicate is  $Clipped(s, f, t)$ , which says that between times  $s$  and  $t$  an event occurs which terminates the process  $f$ . The predicates just introduced are related by axioms, of which we shall see a glimpse below. With this notation, and using  $?\varphi(x)$  **succeeds**

<sup>9</sup>This section can be skipped by readers who have never seen the Event Calculus before.

to abbreviate: ‘make it the case in our discourse model that  $\varphi(x)$ ’,<sup>10</sup> we can write the two update instructions involved in comprehending the discourse as:

(10)  $?Happens(e_1, t), t < now, Happens(e', t')$  **succeeds**

(11)  $?Happens(e_2, s), s < now, Happens(e'', t'')$  **succeeds**

Here  $e'$  and  $e''$  are variables for event types in the discourse context which have to be found by substitution or, more precisely, *unification*. These two update instructions have to be executed so that  $e'' = e_1$  and  $s < t''$ . If ‘Max fell’ is the first sentence of the discourse, we may disregard  $e'$ .<sup>11</sup> In order to formulate the causal knowledge relevant to the execution of these instructions, we introduce a process  $f$  (*falling*) corresponding to the event  $e_1 = fall(m)$ , where  $f$ ,  $e_1$  and  $e_2$  are related by the following statements:

(12)  $HoldsAt(f, t) \rightarrow Happens(e_1, t)$

(13)  $Initiates(e_2, f, s)$

The system processing the discourse will first satisfy the update request corresponding to ‘Max fell’ by locating the event  $e_1$  in the past of the moment of speech. The second sentence, ‘John pushed him’, is represented by the request (11) which contains the variable  $e''$ . The system will try to satisfy the goal by reducing it using relevant causal knowledge. Applying (12) and unifying<sup>12</sup>  $e'' = e_1 = fall(m)$ , the second request (11) is reduced to:

(14)  $?Happens(e_2, s), s < now, Happens(e_1, t''), HoldsAt(f, t'')$  **succeeds**

The system now applies a general causal principle, known as *inertia*, which says that, if an event  $e$  has kicked off a process  $f$  at time  $t$ , and nothing happens to terminate the process between  $t$  and  $t'$ , then  $f$  is still going on at  $t'$ . This principle rules out spontaneous changes, that is changes which are not caused by occurrences of events. Inertia can be formulated as the following axiom:

(15)  $Happens(e, t) \wedge Initiates(e, f, t) \wedge t < t' \wedge \neg Clipped(t, f, t') \rightarrow HoldsAt(f, t')$

Using this axiom, the request (14) is further reduced to:

(16)  $?Happens(e_2, s), s < now, Happens(e_1, t''), HoldsAt(f, t''), Initiates(e, f, t), t < t', \neg Clipped(t, f, t'')$  **succeeds**

<sup>10</sup>This notation derives from logic programming. By itself,  $? \varphi(x)$  denotes a *goal* or *query*, a request for a value  $a$  of  $x$  such that  $\varphi(a)$  is true. The answer may be negative, if the database against which  $\varphi(x)$  is checked contains no such individual. By  $? \varphi(x)$  **succeeds** we mean that in such cases the database must be updated with an  $a$  making  $\varphi$  true. These instructions or requests for updates are also known as *integrity constraints*.

<sup>11</sup>Here we regard context as provided by the preceding discourse, but one may conceive of ‘forward-looking’ notions of context as well.

<sup>12</sup>This form of unification will be important in our discussion of the ‘binding problem’ for language.

Using (13) and unifying  $e = e_2 = \text{push}(j, m)$  and  $s = t$  we reduce this request to:

$$(17) \text{ ?Happens}(e_2, s), s < \text{now}, \text{Happens}(e_1, t''), s < t'', \neg \text{Clipped}(s, f, t'') \text{ succeeds}$$

This is a definite update request which almost says that *push* precedes *fall*, except for the formula  $\neg \text{Clipped}(s, f, t'')$ , which expresses that *f* has not been terminated between *s* and *t''*. If *f* were terminated between *s* and *t''*, we would have a situation as in:

$$(18) \text{ Max fell. John pushed him a second time and Max fell all the way to the bottom of the pit.}$$

Since we have no positive information to this effect, we may assume  $\neg \text{Clipped}(s, f, t'')$ . This form of argument is also known as *closed world reasoning*: ‘assume all those propositions to be false which you have no reason to assume to be true’. Closed world reasoning is essential to planning, and to discourse comprehension, as it allows one to discount events which are logically possible but in practice irrelevant. The final update request is thus:

$$(19) \text{ ?Happens}(e_2, s), s < \text{now}, \text{Happens}(e_1, t''), s < t'' \text{ succeeds}$$

which is the instruction to update the discourse model with the past events  $e_1$  and  $e_2$  such that  $e_2$  precedes  $e_1$ .

Just as plans may have to be revised in mid-execution (for instance, if it turns out there is not sufficient oil to produce the projected number of pancakes), discourse models may have to be recomputed when additional information is provided. Suppose the discourse does not stop after ‘John pushed him’ but, instead, continues:

$$(20) \text{ Max fell. John pushed him, or rather what was left of him, over the edge.}$$

One obvious interpretation is that now  $e_2 = \text{push}(j, m)$  comes after  $e_1 = \text{fall}(m)$ . This is the result of a recomputation, since after the first ‘him’ the hearer may have inferred that  $e_2$  precedes  $e_1$ . Let us give a brief, informal sketch of this recomputation. The phrase ‘or rather what was left of him’ suggests Max is now *dead*, therefore the update request corresponding to the second sentence is something like:

$$(21) \text{ ?Happens}(e_2, s), s < \text{now}, \text{HoldsAt}(\text{dead}(m), s), \text{Happens}(e'', t'') \text{ succeeds}$$

perhaps together with a requirement to the effect that the entire pushing event occurs while *dead*(*m*) obtains. It now seems reasonable to assume that, at the start of *falling* (the process denoted by *f*), Max is still *alive*. Unifying  $e'' = e_1$  and applying property (12), the request reduces to finding instants  $s, t''$  such that:

$$(22) \text{ ?Happens}(e_2, s), s < \text{now}, \text{HoldsAt}(\text{dead}(m), s), \text{HoldsAt}(\text{alive}(m), t''), \text{Happens}(e_1, t'') \text{ succeeds}$$

can be satisfied. Since *alive* always precedes *dead* and not conversely, it follows that we must have that  $e_1 = \textit{fall}$  precedes  $e_2 = \textit{push}$ .

In summary, what we have outlined here is a computational mechanism for determining event structure from discourse, based on planning. Temporal expressions are hypothesized to determine requests to be satisfied by an update of the current discourse model. Processing these requests involves unification, search through semantic memory, as well as setting up temporary structures in working memory.

### 3.2.2 Computing event structures for (PS, Imp) combinations

Similar arguments apply to the French examples with which we started this section:

- (7) a. Il faisait chaud. Jean ôta sa veste. (Imp, PS)  
*It was hot. Jean took off his sweater.*

Intuitively, this narrative determines an event structure in which *hot* acts as a background which is true all the time, and the foregrounded event (taking off one's sweater) is located within this background. One arrives at this structure by means of the following argument. World knowledge contains no causal link to the effect that taking off one's sweater changes the temperature. The goal corresponding to the first sentence dictates that it is hot at some time  $t$  before *now*. By the principle of inertia, the state *hot* must either hold initially (at the beginning of the narrative) or have been initiated. The latter requires the occurrence of an initiating event, which is however not given by the discourse. Therefore, *hot* holds initially. Similarly, no terminating event is mentioned, so *hot* extends indefinitely, and it follows that the event described by the second sentence must be positioned inside *hot*.

The second example dates back to the bygone days when speeding cars were stopped by the police instead of being photographed:

- (7) b. Jean attrapa une contravention. Il roulait trop vite. (PS, Imp)  
*Jean got a ticket. He was driving too fast.*

It is given that the event of getting a ticket occurred sometime in the past, and it is also given that the fluent *speeding* was true sometime in the past. Hence, it holds initially or has been initiated. We have to determine the relative position of event and fluent. World knowledge yields that getting a *ticket* terminates, but does not initiate *speeding*. Because this is the only event mentioned, *speeding* holds from the beginning of discourse, and is not reinitiated once it has been terminated.

In the third example, the same order of the tenses yields a different event order, guided by the application of causal knowledge:

- (7) c. Jean appuya sur l'interrupteur. La lumière l'éblouissait. (PS, Imp).  
*Jean pushed the button. The light blinded him.*

One (occurrence of an) action is mentioned, pushing the light button, which has the causal effect of initiating the light being on when its current state is off. No terminating event is mentioned, therefore the light remains on. It also follows that the light must be off for some time prior to being switched on, and that it must be off at the beginning of discourse. The definite article in ‘*La lumière*’ leads to a search for an antecedently introduced light, which successfully terminates after unification with the light introduced in the first sentence. As a consequence, it is *this* light which is too bright.

### 3.2.3 *Deviant verb tenses and ADHD*

In cognitive terms, planning is part of ‘executive function’, an umbrella term for processes responsible for higher-level action control which are necessary for maintaining a goal and achieving it in possibly adverse circumstances. Executive function comprises maintaining a goal, planning, inhibition, coordination and control of action sequences. Since we have postulated that tense processing involves planning toward goal, we see that several components of executive function are involved in comprehension and production of tense and aspect. A corollary is that failures of executive function can show up in deviant use of tense and aspect and in impairments in processing temporal discourse, for instance in ASD (Autistic Spectrum Disorder), ADHD (Attention Deficit Hyperactivity Disorder), and schizophrenia. Of particular interest here will be children with ADHD, a disorder that is characterised by persistent and developmentally inappropriate levels of inattention, impulsivity and hyperactivity. About 2% of children (mainly boys) are severely affected; 3–6% suffer from less severe symptoms.<sup>13</sup> It has been hypothesized to be an executive function disorder, and indeed children with ADHD score significantly lower on a number of standard tests measuring components of executive function, such as planning, inhibition, and self-monitoring. The precise pattern of executive deficits in ADHD is not yet known, and it is not yet determined whether there is a single executive deficit that explains most of the symptoms. Below we will investigate linguistic consequences of the hypothesis that goal maintenance is affected in ADHD, evidence for which can be found in [Shue and Douglas, 1992; Pennington and Ozonoff, 1996]. For instance, it is known that children with ADHD have trouble with retelling a story, a task that involves presenting information so that it is organized, (temporally) coherent, and adjusted to the needs of the listener. The ability to attend to these requirements presupposes that one is able to retain goals in working memory while planning the necessary steps and monitoring their execution. This ability requires executive function as defined above [Purvis and Tannock, 1997], and is known to be compromised in ADHD. On difficulties with in maintaining goals in working memory, see [Geurts, 2003].

Given that goal maintenance in working memory is compromised in children with ADHD, together with the proposal that such maintenance is necessary to allow computation of event structures (i.e. tense processing) we are led to the

---

<sup>13</sup>These are figures for the Netherlands, supplied by the *Gezondheidsraad*.



following suggestion [van Lambalgen *et al.*, 2008].

Recall that update requests, that is the goals to be satisfied, corresponding to a VP's tense and aspect, consist of two components:

1. location of an event in time;
2. meshing the event with other events.

If a child has trouble maintaining a goal in working memory, this may lead to a simplified representation of that goal. In the case of verb tenses, the most probable simplification is of 'location of event in time' (never mind the meshing with other events), since this involves the least processing (search through semantic memory and unification). This simplification affects both comprehension and production, the case of interest here. Indeed, in producing speech which is attuned to the needs of the listener, the speaker may construct a discourse model of his own utterances, to determine whether it is sufficiently unambiguous. Slightly more formally, our main hypothesis is:

A speaker with ADHD simplifies the goals corresponding to tenses at the expense of the hearer.

We list here a number of ways in which these goals can be simplified. An extreme form of complexity reduction is not to use tensed forms at all. For example, in frog-story experiment on ADHD narration, we saw discourses like this

En hij is vroeg op. En wat ziet die daar? Kikker verdwenen.  
*And he is up early. And what does he see there? Frog gone.* [7yrs, ADHD]

The difference between control and ADHD children was quite striking: only 2.9% of controls used tenseless utterances in their narratives, whereas 19.2% of the ADHD children did so.

A second way in which the child with ADHD can ease the burden on himself while increasing that of the hearer, is using reported speech ('quotes', 'direct speech' only). Here's an example of the contrast between the two groups: two ways of narrating the same scene, that in which the boy climbs into a tree and looks into a hole, whereupon an owl appears in the hole, and knocks the boy out.

a. En die jongen ging zoeken in de boom. En toen zag die een uil. En toen valt 'ie van de boom.

*And that boy started looking in the tree. And then he saw an owl. And then he falls from the tree.* [8yrs, CG] <sup>14</sup>

b. 'Oh nee, ik val!' 'Hellup!' 'Ga weg, stomme uil, ga weg!'  
*'Oh no, I'm falling!' 'Help!' 'Go away, stupid owl, go away!'* [9yrs, ADHD]

<sup>14</sup>The child makes a mistake in combining present tense 'valt', which could be interpreted as a narrative present heightening the tension, with the adverbial 'en toen', which needs a past tense.

There are several other ways the child with ADHD can reduce the complexity of its goals, e.g. reducing the number of context setting elements, or avoiding the perfects (which are computationally intensive). We may now take a more global view, and look at the children who apply one or more complexity-reducing strategies. For example, a child may use up all his computational resources by avoiding direct speech, thereby producing, say, more erratic shifts in the perfect. Both in case of excessive use of direct speech and of erratic tense shifts the hearer must work hard to construct a coherent story, even though he may not understand why he has to work so hard. Thus, taking the point of view of the hearer, what is necessary is a general definition of complexity-reducing strategy, incorporating the more specific strategies discussed above. Motivated by the analyses given above, we define the overall complexity-reducing strategy of a child as consisting of three components: tenseless utterances, direct speech, avoidance of the perfect. For the precise definition of ‘strategy’ we refer the reader to [van Lambalgen *et al.*, 2008], here we state only the main result, that children with ADHD use a strategy with the aim of reducing tense complexity significantly more often. This may explain the sense of unease a hearer feels when listening to such narrative.

Before we close this section, one last word on methodology. The predictions concerning the use (or non-use) of verb tenses in ADHD were derived from a formal model [van Lambalgen and Hamm, 2004] of tense production and comprehension involving satisfaction of complex goals, together with neuropsychological evidence indicating difficulties with goal maintenance and/or planning toward that goal. The formal model is responsible for the specificity of the predictions. Without the formal model, but equipped only with, say, Trabasso and Stein’s general characterisation of narrative as governed by a hierarchy of goals [Trabasso and Stein, 1994], one expects some breakdown in the coherence of story-telling in ADHD, as was indeed found by Purvis and Tannock [Purvis and Tannock, 1997]. The formal model allows one to be more specific about the computational cost of the devices used to ensure discourse coherence. The model thus acts as a searchlight that allows one to see phenomena one would not have thought of otherwise.

#### 4 THE BINDING PROBLEM FOR SEMANTICS

The goal of a theory of language is to deliver analyses at each of Marr’s levels, and to bridge them in a perspicuous manner. One way of achieving this is to define a notion that acts as a ‘wormhole’ [Hurford, 2003] connecting linguistic structures, algorithms, and neurobiological events. A candidate notion is that of ‘unification’, which has been applied on several occasions in this chapter. Below we provide a broad, neuroscience-oriented framework for the concept of unification.

An influential statement of the ‘binding problem’ for cognitive representations is due to [von der Malsburg, 1981], who regarded the binding approach to brain function as a response to the difficulties encountered by classical connectionist networks. Von der Malsburg 1999 refers to a well-known example by [Rosenblatt, 1962] to illustrate the issue. Consider a network for visual recognition constituted

by four output neurons. Two neurons fire when a specific shape (either a triangle or a square) is presented and the other two fire depending on the shape's position (top or bottom of a rectangular frame). So, if there is a square at the top, the output will be [square, top]. If there is a triangle at the bottom, the output will read [triangle, bottom]. However, if a triangle and a square are presented simultaneously, say, the triangle at the top and the square at the bottom, the output would be [triangle, square, top, bottom], which is also obtained when the triangle is at the bottom and the square at the top. This is an instance of the 'binding problem'. Malsburg writes:

The neural data structure does not provide for a means of binding the proposition top to the proposition triangle, or bottom to square, if that is the correct description. In a typographical system, this could easily be done by rearranging symbols and adding brackets: [(triangle, top),(square, bottom)]. The problem with the code of classical neural networks is that it provides neither for the equivalent of brackets nor for the rearrangement of symbols. This is a fundamental problem with the classical neural network code: it has no flexible means of constructing higher-level symbols by combining more elementary symbols. The difficulty is that simply coactivating the elementary symbols leads to binding ambiguity when more than one composite symbol is to be expressed. [von der Malsburg, 1981, p. 96]<sup>15</sup>

Examples of the binding problem are bistable figures such as Necker's cube and Jastrow's duck-rabbit, where the exact same visual features of the stimulus lead to two incompatible representations, depending on how these features are bound together. Since the availability of different representations essentially depends upon the geometric properties of the figure, rather than upon the constitution of perceptual systems as would be the case, for example, for after images [Marr, 1982, pp. 25-26], bistability requires an explanation at Marr's computational level, where properties of stimuli are described and related to information processing goals. Without a characterization of the geometric properties of the figure, and of the mappings between the figure and the two different entities which it can stand for, there would be no basis upon which to claim that the two representations are mutually exclusive.

There exist analogous cases of structural ambiguity in language:

- (23) a. The woman saw the man with the binoculars.  
 b. Respect remains.

Example (23a) has two alternative syntactic representations, one in which the phrase 'with the binoculars' is a PP attached to the NP 'the man' (the man that was seen by the woman had binoculars), and another in which it modifies the VP

<sup>15</sup>Different solutions to Rosenblatt's problem are possible. See [von der Malsburg, 1999] for a proposal in line with the binding hypothesis and [Riesenhuber and Poggio, 1999] for an alternative approach.

(the woman used binoculars to see the man). Here too the features of the stimulus lead to two interpretations, depending on which attachment option is eventually pursued. These sentences typically result in specific neurophysiological responses, suggesting that syntactic binding is a genuine information processing problem for the brain. Sentence (23b) also has two possible parses, and this has consequences for its meaning: it can either be used as a directive speech act, if ‘respect’ is the verb and ‘remains’ the object noun; or it can be used as an assertion, if ‘respect’ is the object noun and ‘remains’ the verb.

There are some similarities between perceptual bistability in the visual and linguistic domains, such as the fact that in both cases we seem to ‘flip’ between the two incompatible representations. But there is also a deeper analogy between the two: structural ambiguity is defined at the topmost level of analysis in both cases, as [Marr, 1982, pp. 25-26] pointed out. Without an independent characterization it remains unclear why such representations are mutually exclusive in the first place. Extending Marr’s line of argument, we emphasize that the binding problem for semantics is best formulated at the computational level, although attempted solutions are bound to require significant contributions at all levels of analysis, including – perhaps most interestingly – the level of neural implementation [Hagoort, 2005; Hagoort, 2006].

#### ACKNOWLEDGMENTS

We wish to thank Oliver Bott, Travis Choma, Bart Geurts, Fritz Hamm, Karl Magnus Petersson, Keith Stenning, Martin Stokhof, Julia Uddén, Theo Vosse, Roel Willems and an anonymous reviewer for their useful comments on earlier versions of this chapter. We are grateful to the Netherlands Organization for Scientific Research for support under grant 051.04.040.

#### BIBLIOGRAPHY

- [Asher and Lascarides, 2003] N. Asher and A. Lascarides. *Logics of Conversation*. Cambridge University Press, 2003.
- [Berman and Slobin, 1994] R.A. Berman and D.I. Slobin, editors. *Relating Events in Narrative: A Crosslinguistic Developmental Study*. Lawrence Erlbaum Associates, 1994.
- [Bolinger, 1968] D. Bolinger. Judgments of grammaticality. *Lingua*, 21:34–40, 1968.
- [Boole, 1958] G. Boole. *An Investigation of the Laws of Thought*. Dover, 1958.
- [Bunge, 1984] M. Bunge. Philosophical problems in linguistics. *Erkenntnis*, 21:107–173, 1984.
- [Burge, 2005] T. Burge. *Truth, Thought, Reason. Essays on Frege*. Clarendon Press, 2005.
- [Carminati et al., 2000] M.N. Carminati, L. Frazier, and K. Rayner. Bound variables and c-command. *Journal of Semantics*, 19:1–34, 2000.
- [Carruthers, 1996] P. Carruthers. *Language, Thought and Consciousness*. Cambridge University Press, 1996.
- [Carruthers, 2002] P. Carruthers. The cognitive functions of language. *Behavioral and Brain Sciences*, 25:657–725, 2002.
- [Chomsky et al., 2002] N. Chomsky, A. Belletti, and L. Rizzi. *On Nature and Language*. Cambridge University Press, 2002.
- [Chomsky, 1955] N. Chomsky. *The Logical Structure of Linguistic Theory*. MIT Press, 1955.
- [Chomsky, 1957] N. Chomsky. *Syntactic Structures*. De Gruyter-Mouton, 1957.

- [Chomsky, 1965] N. Chomsky. *Aspects of the Theory of Syntax*. MIT Press, 1965.
- [Comrie, 1976] B. Comrie. *Aspect*. Cambridge University Press, 1976.
- [Comrie, 1985] B. Comrie. *Tense*. Cambridge University Press, 1985.
- [Croft and Cruse, 2004] W. Croft and D.A. Cruse. *Cognitive Linguistics*. Cambridge University Press, 2004.
- [Devitt, 2006] M. Devitt. Intuitions in linguistics. *British Journal of Philosophy of Science*, 57:481–513, 2006.
- [Eberle and Kasper, 1989] K. Eberle and W. Kasper. Tenses as anaphora. In *Proceedings of EACL*, 1989.
- [Featherston *et al.*, 2000] S. Featherston, M. Gross, T.F. Münte, and H. Clahsen. Brain potentials in the processing of complex sentences: An ERP study of control and raising constructions. *Journal of Psycholinguistic Research*, 29:141–154, 2000.
- [Ferreira and Patson, 2007] F. Ferreira and N.D. Patson. The ‘good enough’ approach to language comprehension. *Language and Linguistics Compass*, 1:71–83, 2007.
- [Fitch *et al.*, 2005] W.T. Fitch, M.D. Hauser, and N. Chomsky. The evolution of the language faculty: Clarifications and implications. *Cognition*, 97:179–210; discussion 211–225, 2005.
- [Frege, 1980] G. Frege. On sense and reference. In P. Geach and M. Black, editors, *Translation from the philosophical writings of Gottlob Frege*, pages 56–78. Blackwell, 1980.
- [Geurts and van der Slik, 2005] B. Geurts and F. van der Slik. Monotonicity and processing load. *Journal of Semantics*, 22:97–117, 2005.
- [Geurts, 2003] H. Geurts. *Executive functioning profiles in ADHD and HFA*. PhD thesis, Vrije Universiteit Amsterdam, 2003.
- [Greenfield, 1991] P.M. Greenfield. Language, tools and the brain: The ontogeny and phylogeny of hierarchically organized sequential behavior. *Behavioral and brain sciences*, 14:531–595, 1991.
- [Hagoort, 1998] P. Hagoort. The shadows of lexical meaning in patients with semantic impairments. In B. Stemmer and H.A. Whitaker, editors, *Handbook of Neurolinguistics*, pages 235–248. Academic Press, 1998.
- [Hagoort, 2005] P. Hagoort. On Broca, brain, and binding: A new framework. *Trends in Cognitive Sciences*, 9:416–423, 2005.
- [Hagoort, 2006] P. Hagoort. The binding problem for language and its consequences for the neurocognition of comprehension. In N. Pearlmutter E. Gibson, editor, *The Processing and Acquisition of Reference*. MIT Press, 2006.
- [Haslam *et al.*, 2007] C. Haslam, A.J. Wills, S.A. Haslam, J. Kay, R. Baron, and F. McNab. Does maintenance of colour categories rely on language? Evidence to the contrary from a case of semantic dementia. *Brain and Language*, 103:251–263, 2007.
- [Hauser *et al.*, 2002] M.D. Hauser, N. Chomsky, and W.T. Fitch. The faculty of language: What is it, who has it, and how did it evolve? *Science*, 298:1569–1579, 2002.
- [Hintikka, 1999] J. Hintikka. The emperor’s new intuitions. *Journal of Philosophy*, 96:127–147, 1999.
- [Hurford, 2003] J. R. Hurford. The neural basis of predicate-argument structure. *Behavioral and Brain Sciences*, 26:261–316, 2003.
- [Hurford, 2007] J.R. Hurford. *The Origins of Meaning*. Oxford University Press, 2007.
- [Jackendoff, 1987] R. Jackendoff. On beyond zebra: The relation of linguistic and visual information. *Cognition*, 26:89–114, 1987.
- [Jackendoff, 2002] R. Jackendoff. *Foundations of Language: Brain, Meaning, Grammar, Evolution*. Oxford University Press, 2002.
- [Jackendoff, 2003] R. Jackendoff. Précis of Foundations of language: Brain, meaning, grammar, evolution. *Behavioral and Brain Sciences*, 26:651–65; discussion 666–707, 2003.
- [Jackendoff, 2007] R. Jackendoff. A Parallel Architecture perspective on language processing. *Brain Research*, 1146:2–22, 2007.
- [Johnson-Laird, 1980] P.N. Johnson-Laird. Mental models in cognitive science. *Cognitive Science*, 4:71–115, 1980.
- [Kamp and Rohrer, 1985] H. Kamp and C. Rohrer. Temporal reference in french. Manuscript, Stuttgart, 1985.
- [Kohler, 1925] W. Kohler. *The mentality of apes*. Harcourt Brace and World, 1925.
- [Labov, 1996] W. Labov. When intuitions fail. *Papers from the parasession on theory and data in linguistics. Chicago Linguistic Society*, 32:77–106, 1996.

- [Levelt, 1972] W. Levelt. Some psychological aspects of linguistic data. *Linguistische Berichte*, 17:18–30, 1972.
- [Lewis, 1970] D. Lewis. General semantics. *Synthese*, 22:18–67, 1970.
- [Marantz, 2005] A. Marantz. Generative linguistics within the cognitive neuroscience of language. *The Linguistic Review*, 22:429–445, 2005.
- [Marr, 1982] D. Marr. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Freeman, 1982.
- [McCarthy and Prince, 1994] J.J. McCarthy and A. Prince. The emergence of the unmarked: Optimality in prosodic morphology. In *Proceedings of NELS*, volume 24, pages 333–379, 1994.
- [McGonigle *et al.*, 2003] B. McGonigle, M. Chalmers, and A. Dickinson. Concurrent disjoint and reciprocal classification by *cebus apella* in serial ordering tasks: Evidence for hierarchical organization. *Animal Cognition*, 6:185–197, 2003.
- [McKinnon and Osterhout, 1996] R. McKinnon and L. Osterhout. Constraints on movement phenomena in sentence processing: Evidence from event-related brain potentials. *Language and Cognitive Processes*, 11:495–523, 1996.
- [McMillan *et al.*, 2005] C.T. McMillan, R. Clark, P. Moore, C. Devita, and M. Grossman. Neural basis of generalized quantifier comprehension. *Neuropsychologia*, 43:1729–1737, 2005.
- [Osterhout and Holcomb, 1992] L. Osterhout and P. Holcomb. Event-related brain potentials elicited by syntactic anomaly. *Journal of Memory and Language*, 31:785–806, 1992.
- [Papafragou *et al.*, 2008] A. Papafragou, J. Hulbert, and J. Trueswell. Does language guide event perception? Evidence from eye movements. *Cognition*, 108:155–184, 2008.
- [Pennington and Ozonoff, 1996] B.F. Pennington and S. Ozonoff. Executive functions and developmental psychopathology. *Journal of Child Psychology and Psychiatry*, 37:51–87, 1996.
- [Pinker and Jackendoff, 2005] S. Pinker and R. Jackendoff. The faculty of language: What’s special about it? *Cognition*, 95:201–236, 2005.
- [Pulvermuller *et al.*, 2001] F. Pulvermuller, M. Harle, and F. Hummel. Walking or talking? Behavioral and neurophysiological correlates of action verb processing. *Brain and Language*, 78:143–168, 2001.
- [Pulvermuller *et al.*, 2005] F. Pulvermuller, O. Hauk, V.V. Nikulin, and R.J. Ilmoniemi. Functional links between motor and language systems. *European Journal of Neuroscience*, 21:793–797, 2005.
- [Pulvermuller, 2005] F. Pulvermuller. Brain mechanisms linking language and action. *Nature Reviews Neuroscience*, 6:576–582, 2005.
- [Purvis and Tannock, 1997] K.L. Purvis and R. Tannock. Language abilities in children with attention deficit disorder, reading disabilities, and normal controls. *Journal of Abnormal Child Psychology*, 25:133–144, 1997.
- [Quine, 1970] W. V. Quine. Methodological reflections on current linguistic theory. *Synthese*, 21:386–398, 1970.
- [Riesenhuber and Poggio, 1999] M. Riesenhuber and T. Poggio. Are cortical models really bound by the ‘binding problem’? *Neuron*, 24:87–93, 1999.
- [Rosenblatt, 1962] F. Rosenblatt. *Principles of Neurodynamics: Perceptrons and the Theory of Brain Mechanisms*. Spartan Books, 1962.
- [Shue and Douglas, 1992] K.L. Shue and V.I. Douglas. Attention deficit hyperactivity disorder and the frontal lobe syndrome. *Brain and cognition*, 20:104–124, 1992.
- [Singer, 1994] M. Singer. Discourse inference processes. In M.A. Gernsbacher, editor, *Handbook of Psycholinguistics*. Academic Press, 1994.
- [Spivey and Gonzalez-Marquez, 2003] M.J. Spivey and M. Gonzalez-Marquez. Rescuing generative linguistics: Too little, too late? *Behavioral and Brain Sciences*, 26:690–691, 2003.
- [Steedman, 2002] M. Steedman. Plans, affordances, and combinatory grammar. *Linguistics and Philosophy*, 25:723–753, 2002.
- [Stenning and van Lambalgen, 2008] K. Stenning and M. van Lambalgen. *Human Reasoning and Cognitive Science*. MIT Press, 2008.
- [Stenning, 2003] K. Stenning. How did we get here? A question about human cognitive evolution. Frijda Lecture, University of Amsterdam, 2003.
- [Taylor *et al.*, 2008] L.J. Taylor, S. Lev-Ari, and R.A. Zwaan. Inferences about action engage action systems. *Brain and Language*, 107:62–67, 2008.
- [Toni *et al.*, 2008] I. Toni, F.P. de Lange, M.L. Noordzij, and P. Hagoort. Language beyond action. *Journal of Physiology-Paris*, 102:71–79, 2008.

- [Trabasso and Stein, 1994] T. Trabasso and N.L. Stein. Using goal-plan knowledge to merge the past with the present and future in narrating events on line. In M.H. Haith, J.B. Benson, R.J. Roberts, and B.F. Pennington, editors, *The development of future-oriented processes*, pages 323–352. University of Chicago Press, 1994.
- [van Lambalgen and Hamm, 2004] M. van Lambalgen and F. Hamm. *The Proper Treatment of Events*. Blackwell, 2004.
- [van Lambalgen *et al.*, 2008] M. van Lambalgen, C. van Kruistum, and E. Parriger. Discourse processing in attention-deficit hyperactivity disorder (ADHD). *Journal of Logic, Language, and Information*, 17:467–487, 2008.
- [von der Malsburg, 1981] C. von der Malsburg. The correlation theory of brain function. Internal Report 81-82, Department of Neurobiology, Max Planck Institute for Biophysical Chemistry, 1981. Reprinted in: E. Domany, J.L. van Hemmen and K. Schulten (eds.) *Models of neural networks II*, Springer Verlag, 1994.
- [von der Malsburg, 1999] C. von der Malsburg. The what and why of binding: The modeler's perspective. *Neuron*, 24:95–104, 1999.
- [Yngve, 1960] V.H. Yngve. A model and a hypothesis for language structure. *Proceedings of the American Philosophical Society*, 104:444–466, 1960.
- [Zwaan and Radvansky, 1998] R.A. Zwaan and G.A. Radvansky. Situation models in language comprehension and memory. *Psychological Bulletin*, 123:162–185, 1998.