For Maria Piñango

Genesis of a theory of language: From thematic roles (source) to the Parallel Architecture (goal) (Sort of an intellectual memoir)[1]

Ray Jackendoff

The Parallel Architecture is a framework for understanding language that I have been slowly developing over the past 40 years. It is useful to divide its development into two phases. The first phase, from about 1967 to about 1990, was the working out of a theory of semantics, Conceptual Semantics, intended originally to be bolted onto a standard generative grammar. However, as the theory developed and grew – and as generative syntax changed over the years – the two gradually diverged to the point where they were no longer compatible. Thus the second phase, from about 1990 to the present, has been to look at the consequences of Conceptual Semantics for the rest of the grammar, culminating in the Parallel Architecture itself. Up till recently, the main emphasis has been on the consequences for syntax (Simpler Syntax) and for the relation of syntax to semantics. A third phase is now getting underway in which I am beginning to explore the consequences for morphology, phonology, and further phenomena such as sociolinguistic variation. Here I'll talk about the first two phases, with hints at the end about the new directions.

1. Thematic roles

I first got interested in semantic effects in grammar through Jeffrey Gruber's MIT dissertation (Gruber 1965/1970). It was couched in Gruber's own idiosyncratic formalism, which did not resonate with the generative grammar of the time, although related ideas turned up in Lakoff's dissertation (Lakoff 1965/1971) and a few years later in generative semantics and cognitive grammar, (Lakoff 1970, Talmy 1978, Langacker 1987).

Gruber's major insight was that as one moves from one semantic field to another, the same verbs and prepositions keep cropping up, expressing the same functions, relativized to the particular semantic fields in question. The basic datum was the field of spatial location and motion, for which Gruber offered examples like these:

(1) a. The train **was at** the station.
b. i. The train **went from** New York **to** Chicago.
ii. The train **got** to Chicago.
c. The train **stayed at** the station.
d. The engineer **got** the train **to** Chicago.
e. The dispatcher **kept** the train **at** the station.

[1] In spring 2013, Maria Piñango invited me to talk about the Parallel Architecture in her seminar at Yale. The experience was so successful that I thought it might be interesting to write it down. This is the result. I'm grateful to Maria for the invitation and much discussion, and to Peter Culicover, Anastasia Smirnova, and Jenny Audring for many useful comments.

Gruber introduced the term *theme* for the entity in motion or being located. (1a) expresses a theme – the train – being located at a *location*, the station. (1b) expresses motion of a theme from a *source*, New York, to a *goal*, Chicago. (1c) also expresses location of the theme, but over a period of time, not just at a point in time. (1d) expresses an *agent* causing a theme to move to a goal; (1e) expresses an agent causing a theme to stay at a location. One then can semiformally characterize these examples as in (2).

(2) a. BE (Theme: train; Location: station)
b. GO (Theme: train; Source: New York; Goal: Chicago)
c. STAY (Theme: train; Location: station)
d. CAUSE (Agent: engineer; GO (Theme: train; Goal: Chicago))
e. CAUSE (Agent: dispatcher; STAY (Theme: train; Location: station))

Gruber then turned to sentences expressing possession and change of possession, such as (3).

(3) a. i. John **has** the book.
ii. French: Le livre **est à** Jean.
b. i. Sue **gave** the book **to** John.
ii. John **got** the book **from** Sue.
iii. the gift of a book **from** Sue **to** John
c. John **kept** the book.

(3a) expresses possession at a point in time. The French version uses the verb *be* and a locative preposition, paralleling the English (and French) expressions of spatial location. The expressions in (3b) involve change of possession, with the original possessor as the object of *from* and the final possessor as the object of *to*. (3c) expresses possession over a period of time, using the same verb that spatial location uses for causation of location over a period of time. Gruber therefore suggests that the object being possessed or changing possession is a theme, and the possessors are a sort of location, source, and goal:

(4) a. BE (Theme: book; Location: John)
b. GO (Theme: book; Source: Sue; Goal: John)
c. STAY (Theme: book; Location: John)

In order to differentiate spatial location from possession, the basic functions can be annotated with subscripts: $BE_{Spatial}$ vs. $BE_{Possessional}$.

A third field discussed by Gruber concerned adjectival predication (he called this field Identificational; I later called it Ascriptional). The basic paradigm is illustrated by (5).

(5) a. Bob **is** angry.
b. i. Bob **got** angry.
ii. Bob **went from** puzzled **to** angry in five seconds flat.
c. Bob **stayed** angry all afternoon.
d. The noise **got** Bob angry.
e. The noise **kept** Bob angry.

Again there is a common collection of verbs and prepositions, not exactly identical to those in (1) but displaying considerable overlap. This suggests an analysis along the lines of (6).

(6) a. BE_{Ascr} (Theme: Bob; Location: angry)
b. GO_{Ascr} (Theme: Bob; Source: puzzled; Goal: angry)
c. $STAY_{Ascr}$ (Theme: Bob; Location: angry)
d. CAUSE (Agent: noise; GO_{Ascr} (Theme: Bob; Goal: angry))
e. CAUSE (Agent: noise; $STAY_{Ascr}$ (Theme: Bob; Location: angry))

On my reading of Gruber, his analysis leads to a number of important claims. First, various semantic fields have a similar structure in terms of the basic functions BE, GO, STAY, and CAUSE and the thematic relations Theme, Location, Source, and Goal. The strongest claim, that *all* semantic fields can be so structured, probably turns out to be false.[2] Nevertheless, even if only some significant range of fields turns out to be so structured, this is an important fact about the structure of meaning.

Second, verbs and prepositions that are polysemous with respect to semantic field tend to retain their identity in terms of what basic function they express or what thematic role they mark. Given the arbitrariness of the sign, we should not expect total generality or total consistency, but the parallelisms are noteworthy.

Third, in each semantic field, "location" has a different meaning. In the spatial domain, it amounts to being in a physical location. In the possessional domain, "location" is a person, a potential owner, and the theme is the item that is owned. In the ascriptional domain, "location" is a point in a "property space" such as size, color, value, or emotional affect. "Motion" in each domain is change over time from a beginning "location" – the "source" – to a final location – the "goal." (I later expanded the notion of motion to include not just sources and goals but also intermediate points along a path of motion or change.)

Fourth, a verb's meaning cannot be a semantic primitive, since it must include at least a primitive function BE, GO, STAY, CAUSE, plus a semantic field feature. The primitive functions and the field features are like phonological distinctive features, which cannot be used in isolation, but only when bound up in a complex of features that forms a segment. (They are also like quarks, constituents of elementary particles that can never occur in isolation.)

Gruber's analysis was a major advance on the Katz-Fodor semantics of the time (Katz and Fodor 1963), in that it proposed a function-argument-like notation, a principled set of primitives, and a principled semantic basis for the sort of polysemy that crosses semantic fields. Function-argument notation appeared as well in formal semantics; but thematic roles and semantic fields did not. Rather, these emerged again in metaphor theory (Lakoff and Johnson 1980), without citation of Gruber (whose work Lakoff knew, through my dissertation if nothing

[2] I made this claim in Jackendoff 1983, under the rubric of the Thematic Relations Hypothesis. Counterexamples include psych verbs, which have their own set of roles, and various social predicates such as reciprocity, rights, and obligations. See Jackendoff 2007.

else) and eventually in theories of embodied cognition (Lakoff and Johnson 1999; Varela, Rosch, and Thompson 1991).

2. Application of thematic roles to control in nominals

At that time, the theory of control was dominated by Rosenbaum's (1967) Minimal Distance Principle, which explained nicely why the controller of *leave* is the subject in (7a) but the object in (7b): in the latter, *Sally* is closer to *leave* than is *John*. However, there was a well-known exception, *promise*, for which the expected switch of control did not occur, as seen in (7c,d).

(7) a.	John wanted/expected to leave.	[John leaves]
b.	John wanted/expected Sally to leave.	[Sally leaves]
c.	John promised to leave.	[John leaves]
d.	John promised Sally to leave.	[John still leaves]

There was also a further puzzle in the air concerning control in nominals.[3] (8a-d) form a 2x2 matrix: If either the verb or the nominal is switched, control switches; but if both are switched, control remains the same.

(8) a.	John gave Sally an order to leave.	[Sally leaves]
b.	John got from Sally an order to leave.	[John leaves]
c.	John gave Sally a promise to leave.	[John leaves]
d.	John got from Sally a promise to leave.	[Sally leaves]

These pose a serious problem for the minimal distance principle, and indeed for any purely syntactic solution to control: the syntax is essentially the same in all four examples (except for the addition of *from* in (8b,d)), but control shifts wildly back and forth.

One of the great epiphanies of my research life, around 1968 or 1969, was the realization that control in (8) could be characterized in terms of thematic roles. The idea is this: First, both *promise* and *order* (either as verbs or as nouns) involve the transfer of information from one individual to another. With the verbs *promise* and *order*, the source of the information is in subject position, and the goal of the information is in object position.

Second, a promise concerns an action that is to be carried out by the promiser, i.e. the source: you can only make promises about things over which you take responsibility.[4] In contrast, an order concerns an action that is to be carried out by the person receiving the order, i.e. the goal: you can't order someone for yourself to do something. That is, control with *promise* and *order* aligns with the thematic roles of source and goal respectively. This explains why the facts come out as they do in (7c,d).

---

[3] Jackendoff 1972 refers to this as a "well-known old puzzle." Similar examples appear in Chomsky 1968 (48-49), some attributed to Maurice Gross.

[4] Or at least you can't if you express it as *promise to VP*. You can however promise someone that it will rain or that Santa Claus will come. In these cases, promising amounts simply to giving assurance. In *promise to VP* constructions, the promiser is doing something stronger: undertaking an obligation to perform the action in question. See Culicover and Jackendoff 2005.

Third, the NPs containing *promise* and *order* in (8) don't themselves have subjects or objects, as they do in *John's promise to Sally to leave*. However, they do have *implicit* sources and goals, and crucially, these are determined by the thematic roles of the main verb. As seen above in (3b), *give* assigns the role of source to its subject and the role of goal to its indirect object; *get* assigns the role of source to the *from*-phrase and the role of goal to its subject.

For example, consider (8a). *Give* says that *John* is source and *Sally* is goal; hence the implicit source of *order* is *John* and the implicit goal of *order* is *Sally*. In turn, *order* says that its goal controls its complement; hence *Sally* is controller. In (8b), the main verb says that *John* is goal; hence *John* is implicit goal of *order*, hence it is the controller of *leave*. Similarly in (8c,d), except that this time the source rather than the goal of the main verb becomes controller.

This alignment of the thematic roles of the main verb and the nominal in light verb constructions such as (8) does not come for free in the grammar. In Jackendoff 1974 I demonstrated that there is no reasonable way to derive these effects in terms of mapping a more semantically transparent underlying syntactic structure into the observed surface forms. Instead, I proposed a principle of interpretation that accomplishes this alignment as part of the mapping from syntactic structure to meaning.[5]

I derived two important lessons from this analysis. First, the theory of thematic roles is not just a nice descriptive tool to help account for polysemy; it can do some serious grammatical work. To me at the time, the fact that there are two independent uses for thematic roles constituted important evidence for the validity of the analysis. Second, this is a case where semantic structure is essential to an account of the syntactic phenomenon of control. That is, not only does syntax determine semantics, as just about everybody thought at the time,[6] but in addition, independent properties of semantics determine syntax as well. In this particular case, thematic roles have a variety of syntactic realizations – including as subject, object, object of a preposition, and even as zero in nominals. But whatever the syntactic realization, the same semantic generalization holds.[7]

It should be added that it was possible to see early on that thematic roles do not correspond one-to-one with syntactic positions. We have already seen a case in which a role is not expressed at all as a syntactic argument. For a different case, consider transaction verbs such as *buy* and *sell*. In order for something to be an act of buying or selling, two transfers have to

---

[5] A more detailed version of this proposal appears in Culicover and Jackendoff 2005, including treatment of some long-known counterexamples to subject control with *promise*, which turn out to involve meaning even more deeply – not more syntax. There is now experimental work demonstrating that light verb constructions like *give an order* do indeed require special processing, with some evidence that this extra work is not performed over syntactic structures but rather in the course of semantic integration (many refs here).

[6] -- and many people still think. This view has received renewed prominence in the biolinguistic theory of thought (Berwick and Chomsky 2011), which claims that combinatorial thought is possible only by virtue of being derived from linguistic syntax. See section 3.

[7] The semantic generalization is obscured by the fact that *promise* is about the only verb of English with both transitive syntax and subject control. The generalization becomes much clearer when we look at nominals, where the noun *promise* is joined by words like *vow*, *oath*, *offer*, *commitment*, and *pledge*, whose semantic similarity is evident.

take place: an object changes possession from seller to buyer, and (whether expressed as a *for*-phrase or not) money changes possession from buyer to seller. Thus the seller is simultaneously source of the object and goal of the money, and the buyer is goal of the object and source of the money. In other words, certain verbs can assign multiple thematic roles to the same NP (Jackendoff 1972). Later developments such as the theta criterion (Chomsky 1981) essentially ignored these cases or else redefined thematic roles as syntactic "theta-roles" in such a way as to avoid the semantic issues.[8]

A consequence of this analysis is that thematic roles cannot be defined in terms of diacritics on a semantic representation: both the subject and the object of *buy* would have to be marked as both sources and goals, an unhelpful analysis. Rather, thematic roles have to be defined in terms of a structured semantic representation, here one that defines two subevents, each with its own distinctive array of thematic roles. In other words, thematic roles depend on semantic decomposition of verb meanings.

Overall, these analyses were beginning to lend some plausibility to the notion that semantics has its own internal structure, mirrored only partially in syntax.

3. Why should there be parallels across semantic fields?

If multiple semantic fields display parallel structure, encodable in terms of basic functions and semantic roles, the question of explanatory adequacy arises: Why should linguistic meaning be this way? If syntactic and semantic structure were truly homomorphic, it might be plausible to ascribe these parallelisms to syntax. But as we've seen, the semantic generalizations are to some degree independent of their lexical and syntactic realizations, so syntax can't be the determining factor.

A second epiphany, around 1974, was this answer: *The cross-field parallels in basic functions and thematic roles are telling us something important about the structure of thought!* That is, this analysis places us on track to seek one of the holy grails of linguistics, not to mention of philosophy of language: the nature of human thought as revealed through language. We see already that language is not a *mirror* of thought: thought doesn't yield up its secrets so easily. But this sort of analysis offers the exciting hope of showing us how to carefully pare away all the underbrush of linguistic detail to the *reine Vernunft* hiding behind it.[9]

Immediately a further question arose: Could there be any way to verify hypotheses about the structure of thought, other than through linguistic analysis?[10] If you believe thought is possible only by virtue of having language, then the answer is no.

---

[8] The refinement of the theta criterion late in Chomsky 1981 (p. 335) accounts for these cases. However, the form of the theta criterion that took hold in the literature stemmed from the definition earlier in the book (p. 36), which claimed that every NP is assigned exactly one theta role.

[9] To this day I recall my immediate reaction to the realization: "This is the problem that is going to kill me!" I hope I meant that it was a big enough problem to last me the rest of my life, not that I would die trying to solve it!

[10] I could have and perhaps should have, but did not, ask a prior question: Could there be any way to verify hypotheses about the structure of thought through massive crosslinguistic analysis? Almost everything I've done is on English. But I was going for the gold.

However, as a graduate student I had been deeply impressed by Wolfgang Köhler's 1927 classic *The Mentality of Apes*, in which he describes chimpanzees and other animals solving problems that he has set them. For instance, a banana is hung from the ceiling of a chimp's enclosure, too high for the chimp to reach. The chimp jumps repeatedly at the banana without success, and sulks. At some point he notices a box elsewhere in the enclosure, and perhaps he plays with it a while and then sulks some more. Then suddenly, in one smooth motion, he jumps up, grabs the box, drags it over to a position under the banana, climbs on it, and seizes the banana.

Köhler is intent on disproving accounts of animal behavior in terms of conditioning. He argues that if we saw a human behave like this, we would have no hesitation in saying he or she had been thinking and had come up with a creative plan of action. He shows there are however limitations on what problems the chimps can solve. For instance, if the banana is hung even higher, they never figure out that there is a way to stack boxes so they stand stably, and so they are less reliably successful at the two-box problem.[11] Nevertheless, he establishes (at least to my satisfaction) that there is creative thought going on. Moreover, there is a degree of embedding in such plans: the chimp moves the box "in order to" be able to climb on it, which is "in order to" grab the banana, which is in turn "in order to" eat it. From our point of view, what is most important is that it is thought completely without language.

This is directly relevant to the issue of semantics. If a theory of semantics, based on language, is to yield the structure of thought, we want it to converge with a theory of *non*linguistic thought. To the extent that humans and chimpanzees share a common ancestor, we should expect that human cognition shares a great deal with chimpanzee cognition. Evolution does not throw out its successful solutions, it builds on them. So to the extent that we can form hypotheses about what elements must be present in chimpanzee thought in order to explain their behavior, we should expect similar elements to be present in human thought, though perhaps greatly elaborated.

Within the theory of thematic roles, the spatial domain is an obvious candidate for structured thought that we share with apes. Köhler's chimps clearly understand that the box is in a particular location, that it is possible to move it from one place to another, and that it will stay there if you don't move it. The chimps also understand some notion of causation: if I hit the banana up there with a long stick, it will fall down and I can eat it.

Do chimps' brains encode these notions the same way humans' brains do? Well, we don't know. But a good working hypothesis is that they do, and we can be prepared to back off this strongest hypothesis if evidence arises to the contrary. We might then want to ask whether the extension of the system of thematic relations to other semantic fields is a purely human elaboration. Answering this, of course, depends in part on what we can find out about how chimpanzees or other primates understand the world.

A second important domain of primate thought is *social* cognition: how monkeys and apes understand their dealings with each other. To get a little ahead of my timeline: In 1983-84

[11] Povinelli 2000 likewise demonstrates a considerable range of cases where his chimps fail in problem-solving at points that we would never notice as posing a difficulty.

I had the good fortune to spend a year at the Center for Advanced Study in the Behavioral Sciences at Stanford, and among my fellow Fellows was a group of primatologists – Dorothy Cheney, Robert Seyfarth, Richard Wrangham, Barbara Smuts, and Tom Struhsaker – working on this very issue. I was especially intrigued by the work that culminated in Cheney and Seyfarth's *How Monkeys See the World* (1990) and *Baboon Metaphysics* (2007). They showed that social cognition is a domain built of notions that are only indirectly perceivable, such as dominance, kinship, mating patterns, alliance, group membership, and precursors of theory of mind. Moreover, everything the monkeys do is (somewhat depressingly) familiar, with obvious counterparts in human social behavior.

Again thinking in evolutionary terms, it seems a reasonable working hypothesis that some version of primate social cognition lies at the foundation of human social and cultural understanding. These notions turn out not to be quite so amenable to analysis in terms of the standard thematic roles and basic functions, suggesting that thematic roles, though hugely general, are not the universal structural principle I once thought. I first broached the topic of social cognition in *Languages of the Mind* (1992) and *Patterns in the Mind* (1994), but I did not seriously attempt to integrate it into a formal semantic theory until *Language, Consciousness, Culture* (2007). By that time, there was a flood of work on evolutionary psychology of primate and human social behavior (e.g. Barkow, Cosmides, and Tooby 1992; Hauser 2006; Tomasello 2000; Povinelli 2000, Byrne and Whiten 1988, de Waal 1996), much of which I was able to incorporate into my analysis.

The point of all this for semantic theory is that we don't have to rely entirely on language for our evidence. Primate cognition gives us an independent fix on our quarry, one that offers the potential for connecting the theory of meaning with the rest of the mind.

Another source of evidence for the character of nonlinguistic thought comes from infant studies and language acquisition. Consider learning word meanings: since the sound of a word gives no clue as to its meaning, the meanings children settle on have to be concepts that are accessible to them prelinguistically. Here are some of the kinds of things children learning language are sensitive to (see Carey 2009 for a much more extended list):

- At 17 months, they know that proper names can be given to dolls but not to boxes (Katz, Baker, and Macnamara 1974)
- At 13 months, they know that an action consisting of one person acting on another ought to be distinguished grammatically from an action of two persons doing the same thing independently (Hirsh-Pasek and Golinkoff 1996).
- At 2 years, they know that shape is more relevant for object categorization than color, size, or texture; they also know that animates but not inanimates change shape, so shape is less reliable for animates. (Landau, Smith, and Jones 1988)
- At 28 months, blind children figure out that *see* denotes a form of perception; they take it to be roughly equivalent to *feel* (Landau and Gleitman 1985)
- At 5 months, they have a notion of number (or individuation) sufficient to distinguish one object from two and two from three (Wynn 1992).

- At 5 months, they can interpret the actions of animate and "animatoid" (self-propelled) objects as goal-directed, independent of where the goal is in space (Luo and Baillargeon 2005).
- At 12 months, having observed an animatoid object reaching a goal by detouring around an obstacle, they are surprised when the obstacle is removed and the object still makes the detour instead of going directly – as though they apparently understand some principles of goal-directed behavior (Csibra et al. 2003).
- At 13.5 months, they can attribute to an actor a disposition to perform an action on a novel object, based on its spatial context and its affordances for motion (Song, Baillargeon, and Fisher 2005).

In short, prelinguistic children and children in the very earliest stages of learning language apparently have concepts that structure both their understanding of events they observe and their attempts to assign meanings to words. Again, this gives us an independent fix on a theory of semantics: it should have at its foundations the kinds of concepts that prelinguistic children exhibit.

In view of these nonlinguistic foundations of meaning, I felt it appropriate to use the term *conceptual structure* for the level of representation in which the meanings of sentences are encoded. Accordingly, the first page of *Semantics and Cognition* (1983) declares boldly that "to study semantics of natural language *is* to study cognitive psychology."

4. A detour into music

About the same time as my exploration of conceptual structure was getting underway, I was drawn into a project with the composer Fred Lerdahl on a generative grammar of tonal music. The project was inspired by Leonard Bernstein's Norton Lectures at Harvard in 1973 (Bernstein 1976), in which he proposed a universal grammar for music paralleling the universal grammar for language. Bernstein's lectures led to the formation of a discussion group at MIT on music, language, and aesthetics, which met monthly for about three years (1974-76), and Lerdahl and I began developing our approach in the heady but nurturing context of that group.

Our goal was to ask the same questions about music that generative grammar asks about language: what is the character of the knowledge that listeners experienced in a music genre use to comprehend a piece of music – to hear it as more than a sequence of notes? And where does that knowledge come from? We deliberately did not set out to model our grammar on the structure of language (as Bernstein had); rather we wished the structure of music to speak for itself. Our starting point was Heinrich Schenker's (1935) notion of analysis in terms of a hierarchical "reduction" of a piece of music to successively more structurally prominent elements, such that the final, most prominent element remaining would be a tonic chord. In some ways a Schenkerian analysis can be construed as reflecting the output of a phrase structure grammar, and that is how we began, translating Schenkerian graphs into tree structures that we called *prolongational reductions*.

We discovered very early, however, that the standard generative approach of the time ("generate all and only the grammatical sentences/pieces of music") was not feasible for music.

Any generative system we could develop yielded too many different ways to generate the same piece, most of which had grossly incorrect structures from an intuitive point of view. So we improvised a system in which well-formedness conditions of a familiar generative sort were supplemented by "preference rules," which choose the preferred structure for a piece of music out of the many possible analyses.[12]

It turns out that the preferred structures in prolongational reduction depend very heavily on rhythm: structurally prominent notes or chords tend to (but do not categorically) fall in rhythmically prominent positions. In turn, rhythmic prominence turns out to be a function of two independent structures, *grouping structure* and *metrical structure*, each of which has its own well-formedness conditions and preference rules, and which are loosely linked by further preference rules. Finally, in order to complete the link between rhythm and prolongational reduction, we had to develop a fourth structure, *time-span reduction*, which mediates between rhythmic and harmonic prominence.

The upshot was a theory quite different from classical generative grammar, whose only generative component is syntax, and in which phonology and semantics are "interpretive," that is, derived from the syntax. Our musical grammar required these four independent generative structures, each with its own preferences for optimality. The four structures were furthermore linked with each other by preference rules that specified optimal correspondences. In present terms, these linking rules constitute interfaces among the various structures: rules that create correspondences among diverse forms of mental representation.

This organization of musical grammar became a precedent for the Parallel Architecture in language, in which phonology and semantics are independent from syntax. As I've described in previous sections, Conceptual Semantics was by this time beginning to assert its independence from syntax. But simultaneously, developments in phonology were leading to similar conclusions, for example in metrical theory (Liberman and Prince 1977), autosegmental phonology (Goldsmith 1984), and the study of prosodic contours (Selkirk 1984). The fact that this parallel organization was absolutely necessary for music provided an argument from *outside* linguistics that mental faculties could be organized this way, and it made the Parallel Architecture feel a good deal less strange by the time I formulated it explicitly in the middle 1990s. (And Jerry Sadock was on to some of the same issues by then: Sadock 1991.)

The musical grammar also had a direct effect on the semantic theory, in the mechanism of preference rules. One of the big problems in a theory of meaning is that it's impossible to define almost any word in terms of necessary and sufficient conditions – a situation that led Jerry Fodor (1975) to the bizarre position that the meanings of monomorphemic words such as *telephone* are unstructured and innate (a position he seems to hold to this day). What Fodor could not imagine was a way of combining conditions such that they are not necessary and sufficient.

[12] Notice how this is a precursor of Optimality Theory, in which a highly unselective generative device, GEN, creates candidate analyses, which are then filtered by violable constraints. Alan Prince, one of the originators of OT, was of course familiar with Lerdahl's and my work on music, particularly through the overlaps between the theories of musical and linguistic rhythm.

Lerdahl and I found that the conditions we needed to state for many aspects of musical structure also could not be stated in categorical, necessary-and-sufficient terms. Rather, various conditions could be in competition, one collection of conditions favoring one structure and another collection favoring another structure. At the time this situation seemed very curious, and on a hunch I asked George Miller if it looked at all familiar to him. He suggested that I go look at the gestalt psychologists of the 1920s-1940s. And sure enough, a 1923 paper by Max Wertheimer demonstrated exactly the sorts of effects we had found: sets of conditions (gestalt principles) that affect interpretation, but that can be overridden by one or more other conditions. Each of the conditions, in certain circumstances, can be sufficient to sway an analysis, but in other circumstances it is not necessary and can even be violated. Wertheimer's paper was primarily concerned with conditions in visual perception, but he also mentioned parallels in other perceptual domains, including music. So our notion of preference rules had a precedent after all.

Preference rules also offered one part of the answer to Fodor's despair in lexical semantics. One of the reasons it is impossible to formulate standard definitions for words is that you can find exceptions to just about any condition you think of, as illustrated by Wittgenstein's (1953) famous rant about the meaning of *game*. The solution is to formulate the conditions as preference rules, so that they allow exceptions, if enough of the other conditions are satisfied. This innovation was crucial in putting Conceptual Semantics on a firmer footing, and it played an important role in *Semantics and Cognition* (1983). Even though preference rules are anathema to any sort of semantic theory founded on standard logic, they make perfect sense in a theory that seeks to characterize meanings in the mind. They are a mode of mental computation that can be observed in many different perceptual and cognitive domains, so there is no reason to exclude them from an account of language.

5. How can we talk about what we see?

A further nonlinguistic fix on the character of semantic theory was inspired by a little paper entitled "How do we talk about what we see?", written by my good friend John Macnamara in 1978 but never published. Macnamara's argument was that the visual system has to generate mental representations that create our experience of seeing the world, and that in order for us to talk about what we see, there has to be some sort of translation or conversion from the mental representations arising from the visual system into a form of mental representation usable by the linguistic system. A similar point was made by George Miller and Philip Johnson-Laird in their *Language and Perception* (1976), where they were trying to actually work out a repertoire of predicates usable by both the visual system and language. But Macnamara made the issue very blunt: what could these representations be like, such that one could be converted into the other?

The visual theories in vogue at the time lent themselves to asking this question seriously. Stephen Kosslyn (briefly a colleague of mine) was creating a sensation with his experiments on mental imagery (Kosslyn 1980), suggesting that mental visual images are generated by some sort of analog, space-like representation. A little-discussed aspect of his theory of imagery is that visual images are supported in part by whole-object representations that include even the parts of the imaged object that one cannot "see" at the moment. This is what makes it possible to

perform mental rotation of images on, say, a vertical axis, such that as one part of the imaged object is occluded, new parts come into view.

Even more influential was the work of David Marr (Marr and Nishihara 1978, Marr 1982). Marr set out to develop a theory of object recognition, tracing the computational pathways all the way from the image on the retina to a full representation of object shape. In consonance with linguistic theory, he found it useful to articulate the computation in terms of a series of levels of representation. The first level beyond the retinal image was what he called the *primal sketch*, a representation of "blobs," lines, and edges in the visual field that could be computed on the basis of discontinuities in the retinal image. Marr's next level was the *2½D sketch*, a representation of the visible surfaces and their relative distance from the observer. The final level was the *3D model*, a combinatorial representation of the full shape of the object, independent of the observer's point of view, spelled out in geometric terms. The levels were linked by a variety of computations; perhaps the most impressive part of Marr's account (for me) was the analysis of the multiple sources of depth perception in the computation of the 2½D sketch from the primal sketch.

Here's an example of a Marr 3D model.

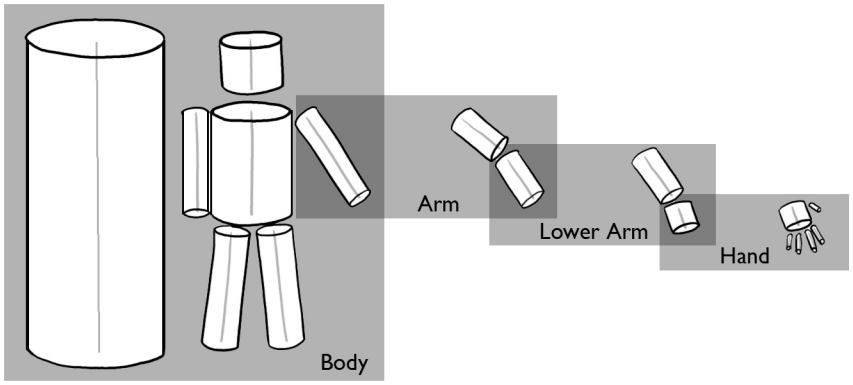

The idea is that at the coarsest layer, a human body is encoded as just a cylinder. When we're interested in the body's overall location, motion, and orientation, this alone is enough. At the next layer, the cylinder is elaborated into parts, each with its own dimensions, orientation, and point of attachment to a larger constituent. This decomposition proceeds recursively, so that for instance one conceptualizes the position and orientation of the fingers in relation to the hand – and then secondarily in relation to the body as a whole. Thus the elaborated structure bears some resemblance to hierarchical phrase structure in language.

Marr made an explicit connection between his approach and linguistic theory. In both approaches, a primary goal is what he called a "computational theory" and what linguists call a "competence theory": an account of the formal structure of the mental representations involved in the cognitive task and of the computations that relate the levels of representation to each other. He stressed also that the computational theory eventually has to be connected to a theory of how these computations are actually carried out and how they are implemented in the brain – a

counterpart to a performance theory in linguistics. Thus the whole philosophy of Marr's approach encouraged a connection with linguistics.

In particular, Marr's approach suggested a direction for answering Macnamara's question of how we talk about what we see. Just as there are computational connections between the levels of representation in visual perception, we should look for computational connections between visual representations and linguistic representations. What would be the appropriate levels to connect?

The obvious connection is between conceptual structure – linguistic meaning – and the 3D model – what might be called visual understanding. When we talk about a cat that we see, we talk about it as a unified single object, independent of our point of view, and independent for instance of the particular shape its tail happens to have at this instant. We don't talk about it as a connected collection of surfaces seen from a particular point of view, much less as a collection of edges and blobs. Hence the 3D model is the appropriate candidate to connect to language. On the linguistic side, we're referring to a cat, and it's conceptual structures, not syntactic or phonological structures, that encode reference.

The relationship between conceptual structures and 3D models can't be established by deriving one from the other, in either direction. First, it is impossible to derive conceptual structure from the 3D model, because conceptual structure contains all sorts of abstract concepts, such as possession, that have no place in visual understanding. There is nothing in the visual appearance of my cat that is relevant to the fact that I own it. Conversely, visual understanding cannot be derived from conceptual structure, because it contains details of shape that are not easily encodable in the algebraic format of conceptual structure. For instance, consider the difference in shape between a duck and a goose. This is encoded naturally in the geometric/typological primitives of the 3D model. But an algebraic feature, say [±long neck], is pretty silly – it's arbitrary and otherwise unmotivated.

Hence the proper solution is to regard the two levels of representation as partners in constructing our understanding of the world. Some elements, such as the notion of physical object, the notion of physical motion, and the notion of physical parts of an object, are shared between the two levels and serve as a conduit (or interface) for aligning the representations with each other. But other sorts of information are proprietary to one representation or the other, such as detailed shape in the 3D model and possession – as well as the type-token distinction! – in conceptual structure. The shared parts form the basis for conceptual structure in the spatial domain. Barbara Landau and I worked out some of the details of how this might work (Landau and Jackendoff 1993), and a great deal of related work on spatial language started developing around this time (e.g. Talmy 1983, Herskovits 1986, Bloom et al. 1996).

Unfortunately, with Marr's untimely death in 1980 (his book was published posthumously), his theory was abandoned by most researchers, partly because of some of its empirical inadequacies, but also partly because it is not part of the tradition in psychology to work with detailed formal theories of mental representation. Shortly after Marr's death, connectionism took hold in psychology, which claimed there is no such thing as mental representation of the sort both linguists and Marr posited. Moreover, about this time began the

revolution in brain imaging, which captured the attention of many researchers: if it was possible to show *where* in the brain something was happening, this seemed more concrete, more like hard science, than formal theories of representation.

Nevertheless, from the standpoint of linguistics, there has to be *some* mental representation, however it is neurally implemented, which encodes our understanding of object shape, orientation, location, and motion, and which can be generalized over categories of objects. I have used the term *spatial structure* for this mental structure, using it as a stand-in for whatever vision researchers eventually come up with. A well-known division is into the ventral and dorsal streams, the so-called "what" and "where" systems (Ungerleider and Mishkin 1982); a considerably more detailed deconstruction is proposed by Landau and Hoffman 2012. (And I'm very disappointed that over the years I've been unable to persuade vision researchers to think in these terms.)

Just as we stepped outside of language per se to justify positing conceptual structure, it's useful to step outside of vision to justify positing something like spatial structure. Consider another way we have of learning about shapes of objects: by feeling them. This haptic modality has an entirely different basis from vision: information comes not from the retinas but from touch and pressure sensors in the skin, plus internal sensors that encode for instance hand configuration. And yet the sense we get of object shape and configuration is not dissimilar: feeling something gives us expectations of what it will look like, and vice versa. This suggests that haptic perception converges with visual perception at the level of spatial structure.

As with the language-spatial interface, there are object characteristics that can only be detected by vision, such as color. But there are other object characteristics that can only be detected by hapsis, such as temperature and hardness. So the overall arrangement of components and interfaces is parallel.

In addition, there is one particular object in the environment for which one has special perceptual information: one's own body. This was brought to my attention by my long-time Brandeis colleague Jim Lackner, who together with Paul Dizio has produced a great deal of research on proprioception (Lackner 1981, Lackner and Dizio 2000). A sense of how the body is configured in space comes from a heterogeneous collection of sensors: muscle spindles, stretch receptors in the joints, touch and pressure sensors in the skin, and the vestibular system in the ears – as well as from visual inspection of the environment and the visible parts of one's body. Again, this source has its own limitations: most notably, you can't feel anyone else's body this way. Nevertheless, all these different sources converge on a unified sense of body position, such that we can match it with the viewed position of someone else's body.

Finally, the understanding of one's own spatial configuration must be combined with the visual and/or haptic understanding of the rest of the environment in order to formulate actions, which eventually get encoded as instructions to the muscles.

Even if we don't know a lot about the representations involved in vision, hapsis, action, and proprioception, from this much we can lay out the overall architecture of these aspects of the mind; and language fits right in as another source of information for the central capacities of

conceptual structure and spatial structure. The difference is that it approaches central cognition from the conceptual structure end. Here is the overall layout; a version of it appears in Jackendoff 1987 (p. 248).

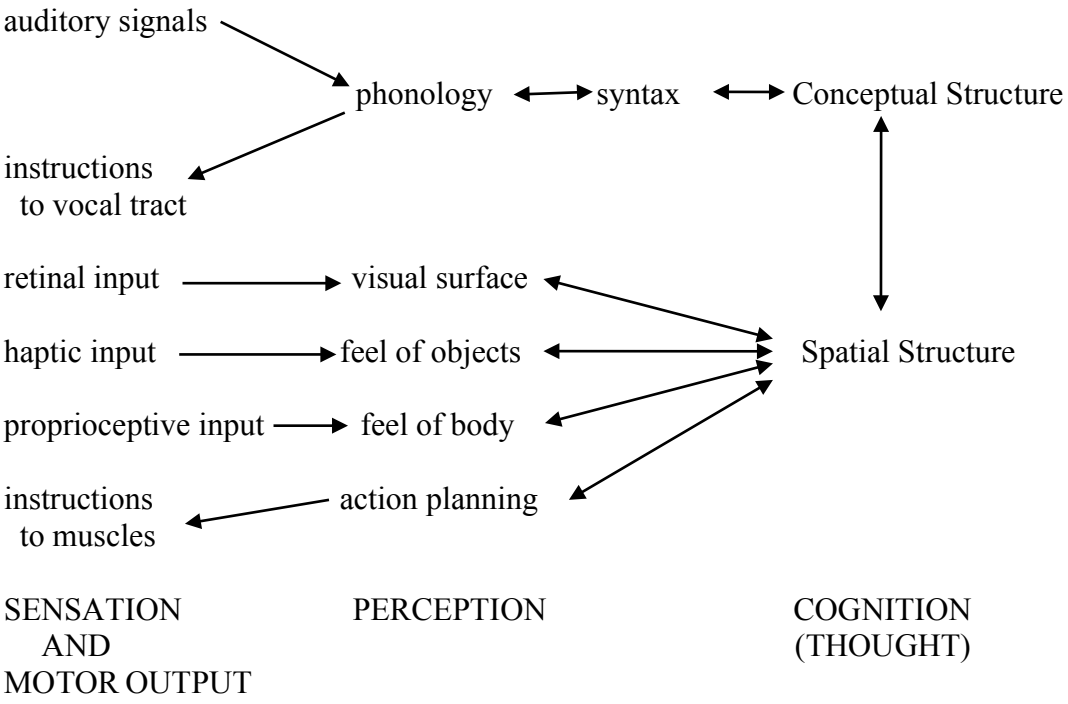

Figure 1. The architecture of the mind

This overview of the architecture of the mind pushes us decisively in the direction of a Parallel Architecture for language, where none of the levels of representation can be *derived* from the others, but they are correlated with each other by way of interfaces.

6. What kinds of entities are there?

In 1976, Jorge Hankamer and Ivan Sag published a paper called "Deep and Surface Anaphora." They were trying to make a point about syntax, but I took from it a somewhat different message. They had examples like this, where a deictic pronoun gets its reference from the visual environment:

(9) [Hankamer tries to stuff a 9-inch ball through a 6-inch hoop.]
Sag: I don't think you can do it.

One typically thinks of demonstrative pronouns as referring to objects in the environment: *Would you pick **that** up, please?* The interest of (9) for me is that here the demonstrative refers to an *action* rather than an object. This led me to a search for more kinds of deictics that denote something other than an object. Here's a list.

(10) a. Type or category:
I'd like one of THOSE. [*pointing to a Porsche*]

b. Location:
Would you put your hat THERE, please? [*pointing*]
There was a fly buzzing around right HERE. [*pointing to empty space*]
The chandelier should hang down to about HERE. [*pointing to empty space*]

c. Trajectory:
The fly went THATAWAY! [*pointing*]

d. Action:
Can you do THIS? [*demonstrating*}
THAT had better never happen in MY house! [*pointing to, say, some random kids smoking pot*]

e. Manner:
Can you walk like THIS? [*demonstrating, say, a Groucho Marx walk*]

f. Distance:
The fish that got away was THIS long. [*demonstrating*]

g. Numerosity:
There were about THIS many people at the party last night. [*gesturing around room or holding up fingers*]

h. Sounds:
Did you hear THAT?

What's the significance of this range of deictic anaphors? The linguistic context tells us what kind of entity to look for in the environment: *this* and *that* in positions that select objects denote objects; in positions that select sounds they denote sounds (10h). In the collocation *do this/that* they denote actions (10d); in the position of a degree word they denote measurements (distances or numerosities) (10f,g). *There* denotes a location; *thataway* denotes a trajectory or direction.

Using a demonstrative this way invites the hearer to find something relevant in the visual environment that fills out the rest of the meaning of the demonstrative. This information has to come through the interface from spatial structure to conceptual structure. In order for it to be of the proper type, say a distance, the visual system must be able to detect distances and encode them in a form that can be linked to conceptual structure. The upshot is that for all these types of entities, the visual system must be prepared to detect them and encode them in spatial structure. In other words, spatial structure has to include a lot more than just objects.

This leads to the converse of Macnamara's question: If we talk about all these kinds of entities as if we see them, how *do* we see all the things we talk about? I take this question to be a challenge for the vision community, which tends to be focused almost exclusively on object perception (with a little side interest in location so layouts of objects can be encoded).

But these examples also have repercussions inside the language system. Most formal treatments of semantics seek to minimize the ontological commitments of the system, if possible restricting the primitive entities to individuals and truth values. Davidson (1967) was taken to be rather daring to propose that the ontology should also include events or actions, and now an

"event variable" is commonplace in formal semantics. What (10) shows is that language acts as though there are also types, places, trajectories, manners, distances, amounts, and sounds – entities one can refer to – and that the visual and auditory system can pick such entities out if the language asks for them.

My conclusion from this is that there is no reason to insist on minimizing the ontology presupposed by natural language semantics. If minimizing were the *only* consideration, well, that would be fine. But the varieties of deictic anaphora show that the mind encodes this considerable range of entities. Since what we're interested in is how the mind codes meaning, we shouldn't hesitate to include all these types in the ontology of the system.

There's a metaphysical point here. Are there *really* places in the world? Distances? Manners? To answer this question, we need to return to the lesson of the gestalt psychologists: our minds *construct* the world of our experience. Consider an ambiguous visual figure like the duck-rabbit. "Out there in the world," there's only a collection of lines on a page or screen. It has two interpretations only because the viewer puts them there. Similarly, for semantics, it doesn't matter whether there are *really* all these entities in the world. What matters is that we all put them into our experience and understanding of the world, and once they are there, we can refer to them. (Some people have suggested to me that this view resembles Kant's; I leave it to someone else to sort out the similarities and differences.)

Yet another kind of deictic anaphora denotes a sentence or proposition:

(11) He really said THAT?

So we talk as though there really are sentences in the world. However, from a more scientific/psychological point of view, a sentence is a *sentence*, rather than a bunch of sound waves, only because there is a person speaking it and perhaps a person hearing it, both of whom connect the sound waves with a concept or idea. That is, a sentence too is a construction in the mind of the speaker and hearer, but it is part of their experienced reality.

Similarly, the truth-values of sentences are not out there in the world. Since truth and falsity are properties of sentences, they too have to be grounded in the minds of language users. The fundamental datum is not the truth of a sentence, an absolute objective fact, but rather a language user's *judgment* that a sentence is true, based on his or her understanding of the relevant situation.

In other words, we're now firmly and irrevocably committed to an internalist theory of meaning. This is perfectly consistent with the internalist view of syntax taken by generative grammar, and quite distinct from much of the tradition in philosophy of language and formal semantics. However, at the same time we have diverged from standard generative grammar in taking semantics to have its own independent generative capacity rather than being derived from syntax, and in seeing semantics as connected directly to the capacities of cognition and perception.

7. Semantic argument structure

These considerations lead fairly naturally to a notation for conceptual structure. For instance, let us return to the Gruber-type examples of (1). (12) shows a labeled bracket notation not unlike predicate logic, and an equivalent but visually more perspicuous tree notation (which I introduced in Jackendoff 2002, following Pinker 1989 and Levelt 1989). The notations such as BE in capitals are to be considered stand-ins for composites of more primitive semantic features.

(12) a. The train was at the station. (= (1a))[13]
[$_{\text{State}}$ BE ([$_{\text{Object}}$ TRAIN], [$_{\text{Place}}$ AT ([$_{\text{Object}}$ STATION])]

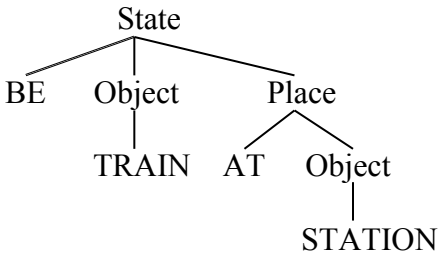

b. The train went to Chicago. (=(1b))
[$_{\text{Event}}$ GO ([$_{\text{Object}}$ TRAIN], [$_{\text{Path}}$ TO ([$_{\text{Object}}$ STATION])]

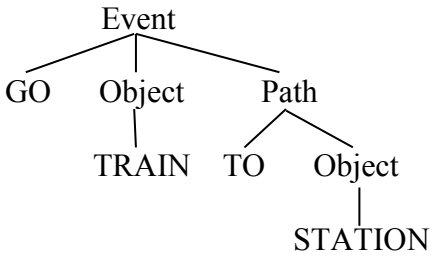

The two notations are related in exactly the way labeled bracketing and tree structure notations are related in syntax. In the labeled bracketing, the ontological category of a constituent is indicated as a subscript on the left bracket; a conceptual function is notated at the left side of a constituent, and its argument(s) are constituents within parentheses. In the tree notation, the ontological category is a node at the top of a constituent; the conceptual function – the head of the constituent – is a left daughter of the category, with its arguments, if any, to its right.

In these notations, thematic roles are not notated explicitly. Rather, theme is the first argument of BE and GO, location is the argument of AT, and goal is the argument of TO. That is, thematic roles have structural definitions, as desired from the considerations in section 2.

[13] This tree is drawn conveniently to parallel English syntax. However, it is meant to serve as the notation for any language's expression of location, whether through postpositions, case-markers, or verbal inflection.

These structures begin to give an idea of the canonical mapping between syntax and conceptual structure. The syntactic arguments – in this case the subject and the subcategorized PP – correspond to semantic arguments. For the most part, heads of syntactic phrases correspond to functions in semantic constituents: BE and GO correspond to verbs; and AT and TO correspond to prepositions. Likewise, if TRAIN and STATION are regarded as zero-argument functions (i.e. constants), they correspond directly to the nouns that head their NPs.

To be a bit more precise, it is useful to ask what parts of the conceptual structures in (12) correspond to the words of the sentence. For instance, the verb *be* cannot just correspond to the function BE. BE alone makes no sense – it is a function with no specification of its arguments. Rather, *be* should correspond to the function BE plus its arguments, which are treated as open variables (underlined):

(13) $[_{State}$ BE $([_{Object}$ X$]$, $[_{Place}$ Y$])]$ or in tree notation,

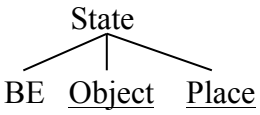

The full conceptual structure (12a) then results from unifying the variables with the conceptual structures of the syntactic arguments of *be*, in this case the subject and the PP. The conceptual structure of the PP in turn is composed by unifying the lexical entry for *at*, (14), with that for *station*.

(14) $[_{Place}$ AT $([_{Object}$ Z$])]$ or in tree notation,

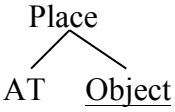

This illustrates the simplest case. A slightly more complex situation can be illustrated with example (1e).

(15) The dispatcher kept the train at the station. (=(1e))
$[_{Event}$ CAUSE $([_{Object}$ DISPATCHER$]$, $[_{Event}$ STAY $([_{Object}$ TRAIN$]$, $[_{Place}$ AT $([_{Object}$ STATION$])])])]$

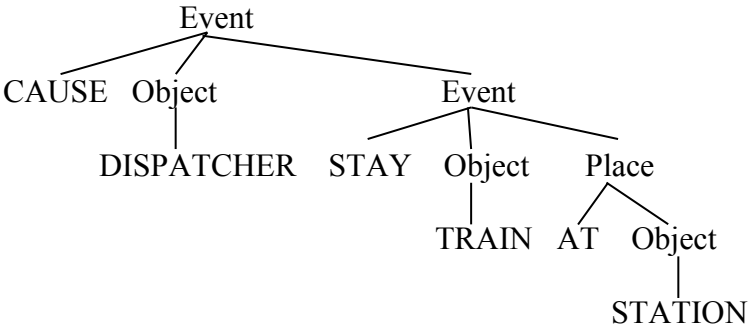

Here, the verb *keep* appears to correspond to two functions in conceptual structure: CAUSE and STAY, a discontinuous piece of the tree. However, if we recognize that semantic argument positions must also be part of the lexical entry, we arrive at (16) for the conceptual structure of this reading of *keep*.

(16)

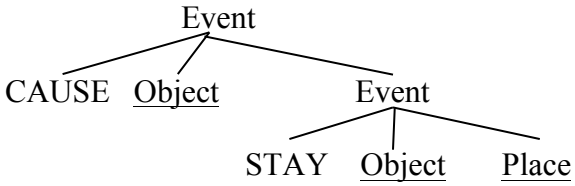

This is now a continuous piece of conceptual structure, with variables corresponding precisely to the three syntactic arguments. Moreover, the syntactic subject corresponds to the highest argument in the semantic tree, which gives us a structural rationale behind the linking hierarchy according to which agents (causers) have priority for subject position.[14]

This is not the only circumstance in which a verb expresses more than one function. Gruber talks about what he calls "incorporation," in which a verb includes a spatial function that could be expressed independently by a preposition. For instance, *enter* can be paraphrased (at least for a first approximation) by *go into*. (17) shows how this is expressed in the formalism.

(17) a. John entered the room = John went into the room

b. Conceptual structure of *go*:

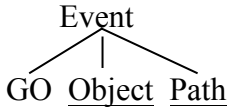

c. Conceptual structure of *into*:

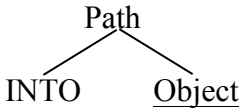

d. Conceptual structure of *enter*:

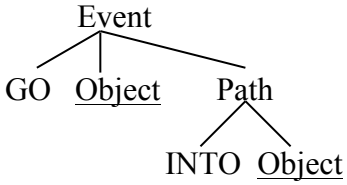

Again this is a connected piece of the semantic tree. The result of "incorporating" INTO with GO is that (17d) has two open arguments, both Objects. Since Objects are always expressed by NPs, *enter* comes out as a transitive verb. Here, the theme is higher in the semantic tree, so it

[14] Notice that this analysis expresses the same insight as the generative semanticists' abstract verb CAUSE (e.g. McCawley 1968; Lakoff 1971). The difference is that the present analysis has no extra verb in syntax, and hence there is no need for a "predicate raising" transformation that collapses the two predicates into a single verb. Rather, the single verb in syntax maps directly into a complex predicate in semantics. That is, the complexity of this verb is localized in the syntax-semantics interface, not in the syntactic derivation.

Jerry Fodor (1970) offers arguments that this semantic analysis (in the generative semantics treatment) is incorrect. For responses, see McCawley 1978, Jackendoff 1983, 2002).

claims the position of subject in syntax. Similar analyses obtain for verbs such as *leave* (= 'go out of/away from'), *cross* (= 'go across'), and transitive *climb* (= 'go to the top of').

Incorporation can go still further. Consider the verb *butter*, which for a first approximation means 'put butter on X'. This can be treated as in (18).

(18) NP butter NP (= 'NP cause butter to go on NP')

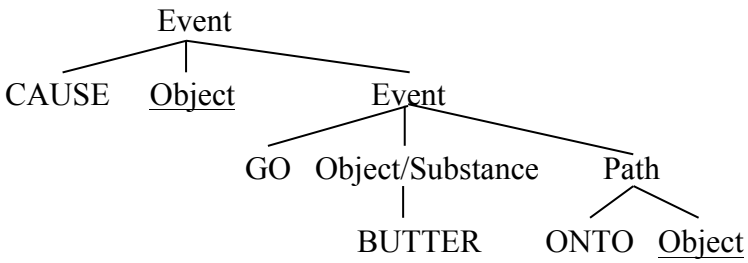

This time, one of the arguments, the theme of GO, is fully instantiated and need not be expressed by anything in the syntax. There remain two unfilled arguments, both objects; hence the verb correctly turns out to be a transitive verb.[15]

A different pattern of incorporation appears in a verb like *bottle* (as in *bottle the wine*). Note that (19) is basically the same semantic structure as (18), but a different argument is filled in: the goal instead of the theme.

(19) NP bottle NP (= 'NP cause NP to go into bottles')

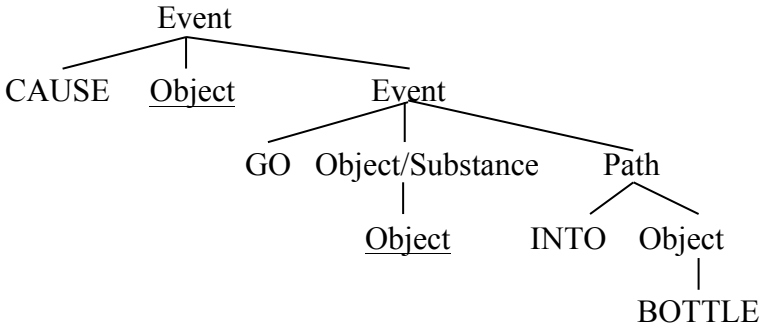

[15] This analysis expresses the same insight as Hale and Keyser's (1993) treatment of these verbs, based on Baker's (1988) notion of *syntactic* incorporation. The difference is that the present analysis has no noun and preposition in syntax, and hence there is no need for a "head to head raising" transformation that collapses the parts into a single verb. Rather, the single verb in syntax maps directly into a complex predicate in semantics. That is, the complexity of this verb is localized in the syntax-semantics interface, not in the syntactic derivation.

The Hale-Keyser account suffers from the fact that the verb *butter* means more than 'put butter on.' You can't butter a dish by putting a stick of butter on it. Buttering involves using butter according to its "proper function," namely spreading it on bread or the like. Also, as we'll see in a moment, buttering doesn't have to involve butter per se. These refinements of meaning can be built into a conceptual structure like (18). They can't be built into a syntactic structure. See Culicover and Jackendoff 2005, 53-56, 102-103.

Hence there are many different ways for a conceptual structure to be linked to a transitive verb, of which (17)-(19) are three. Notice that this analysis sends one off in quite a different direction from Baker's (1988) UTAH (Uniformity of Theta Assignment Hypothesis), according to which every thematic role corresponds to a unique syntactic configuration. Rather, the correspondence of thematic roles to syntactic positions is many-to-many.

A related issue concerns what semantic arguments are obligatorily expressed, and which are optional. For instance, *eat* has transitive and intransitive variants, but the semantically similar *devour* is only transitive.

(20) a. Joe ate/devoured the banana.
b. Joe was eating/*devouring.

However, the semantics of intransitive *eat* still includes an implicit argument: you can't eat without eating *something*. So *eat* and *devour* are both two-argument predicates in conceptual structure, but they differ in syntactic argument structure: *devour* requires its patient argument to be expressed, but *eat* does not.

By contrast, consider *swallow*, which is also optionally transitive.

(21) a. Joe swallowed the banana.
b. Joe swallowed.

Notice that you can swallow without swallowing anything. Hence (21b), unlike (20b), expresses a one-argument predicate, namely performing the bodily function of swallowing. So *swallow* has an optional *semantic* argument – something being swallowed. Thus in the case of *swallow*, the syntactic alternation corresponds directly to the semantic alternation; whereas in the case of *eat*, the syntactic alternation is a consequence of optionality in the syntax-semantics interface.

These considerations (and many more like them) lead to a canonical mapping between conceptual structure and syntactic structure. In particular, the term "argument structure" has to be made more specific: syntactic and semantic argument structure correspond in some respects but diverge in others.

- Syntactic heads correspond to conceptual functions, including complex functions that incorporate connected parts of the semantic tree.
- Arguments of a syntactic head correspond to semantic arguments that instantiate conceptual structure variables of the corresponding conceptual function.
- Arguments higher in the syntactic tree correspond to arguments higher in the semantic tree.
- Certain conceptual arguments of certain verbs need not be expressed in syntax; this is an interface property of particular lexical entries.[16]

And this is approximately where things stood in *Semantics and Cognition* (1983).

[16] In English! In so-called pro-drop languages, there are further principles in the syntax-semantics interface that permit certain arguments to be unexpressed.

8. Mismatches and constructions

But then along came examples like these, which I explored in *Semantic Structures* (1990).

(22) a. Joe buttered the toast with rancid margarine.
b. Joe bottled the wine in attractive bottles.

In these examples, the object of the preposition instantiates the same thematic role as the incorporated material: *rancid margarine* is theme – the substance undergoing change in position, and *attractive bottles* is goal – the place where the wine ends up. A more canonical expression of the same meaning is (23).

(23) a. Joe spread rancid margarine on the toast.
b. Joe put the wine into attractive bottles.

The extra NP in sentences like (22) has to provide additional semantic material beyond what has been incorporated:

(24) *Joe buttered the toast with butter.
*Joe bottled the wine in bottles.

Should this extra phrase be treated as an argument or as an adjunct? From the point of view of syntax, it's an adjunct, since it's completely optional. But from the point of view of semantics, it overlays a semantic argument with supplemental information, so it is sort of an argument.

These constructions further undermine any idea of a one-to-one mapping between syntax and semantics. For one thing, there are two words in the sentence that express the same thematic role: the verb *butter*, implicitly, and the *with*-phrase explicitly. This *with*-phrase of theme is possible with many (all?) verbs that have an incorporated theme. We now are faced with three ways that a theme can be expressed syntactically: as a subject (e.g. with *enter*), as an object (with *keep*), and as a *with*-phrase (with *butter*). This situation thus argues for a flexible interface between syntax and semantics.

Still more problematic is this configuration.[17]

(25) Bill belched his way out of the restaurant.

The phenomenon is productive: it's possible to use any verb of the right semantic type here.

(26) a. Bill drank his way across the country.
b. Bill knitted his way through the conference.

[17] Another epiphany, dating from spring 1987, when I was visiting University of Arizona for a semester.

There are many difficulties here. First, the verb *belch* doesn't usually take a direct object, but (25) has this weird thing *his way* in object position. We can tell it's in object position because the verb can't take its own direct object; evidently *his way* has co-opted object position:

(27) a. Bill drank (*beer) his way across the country.
b. Babe Ruth {homered/*hit home runs} his way into the hearts of America.

Second, *belch* also doesn't select a path expression like *out of the restaurant*. Third, the meaning of the sentence involves Bill *going* out of the restaurant, even though there's no verb of motion. What is *belch* doing, then? It's describing the *manner* in which Bill goes out of the restaurant.

In short, (25)-(26) have a radically non-canonical mapping between syntax and semantics, for which the phrase *his way* in object position is the overt cue. The direct object and the PP are apparently not licensed by the verb; the meaning GO, of which the PP is the path argument, is not expressed in syntax; and the main verb in syntax functions in semantics as a manner or means modifier.

How should this curious situation be accommodated in the grammar? In *Semantic Structures*, I worked through three possibilities. The first was an attempt to derive (25) syntactically from a more canonical structure such as *Bill went out of the restaurant, belching*; this failed miserably. The second was a lexical rule that converted the verb *belch* into an idiom *belch his way PP*, along the lines of lexical rules in LFG and HPSG; this idiom in turn would undergo canonical argument structure realization, resulting in (25).

This second solution at least worked. But it struck me as fishy. Here's why. The point of idioms is that they are supposed to be idiosyncratic and listed in the lexicon. However, the construction in (25)-(26) is productive. One shouldn't have to list every occurrence in the lexicon, and it seems odd that every time one learns a new verb with the appropriate semantics, one also has to store its *his way* counterpart as well.

Well, one might say (and various people *did* say), the lexicon isn't *just* a list of stored forms, it can also contain "potential forms" that can be created on the spot. This proposal was made mostly in the context of morphology in frameworks like LFG and HPSG. For instance, in these frameworks, the passive forms of verbs, with their noncanonical mapping of semantics to syntactic argument structure, are taken to be produced by "lexical rules," then inserted into syntactic structures, where normal principles of argument realization apply. Given the productivity of passive, it seems odd to think of all passive forms being listed in the lexicon. So "lexical rules" might be taken to be rules that produce "lexical items" that are not listed in the lexicon. In effect, "lexical items" might be taken as equivalent to "words."

Even if one feels comfortable with talking about the passive this way, its artificiality becomes more apparent in talking about (25). What's being produced by the putative lexical rule is not a word but a syntactic structure. So it's only in a very attenuated sense of "lexical" that *belch his way PP* is a lexical item: it is neither a word nor is it stored.

Faced with this discomfort, I proposed a third treatment of (25), which I soon discovered was very close to what Adele Goldberg was proposing for the very same phenomenon in the then new and innovative framework of Construction Grammar (Goldberg 1995). The idea is that just as there is a special interface rule for the *with*-phrases in (22), there is an interface rule that directly expresses the mismatch of syntax and semantics in the *way* construction. (28) gives the basic formalism behind the rule; double arrows show the correspondence between syntactic and semantic constituents (*X* in the semantics corresponds to the subject; *pro* in the syntax is a bound pronoun parallel to that in *crane pro's neck*. I leave the treatment of these two pieces aside.)

(28) Semantics: $[_{\text{Event}}$ GO (X, Path); Manner/Means: $[_{\text{Event}}$ F (X)]]

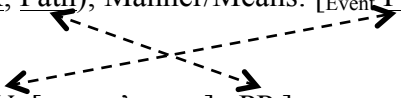

Syntax: $[_{\text{VP}}$ V $[_{\text{NP}}$ pro's way] PP]

The semantics of (28) is identical to that of the more canonical paraphrase, *Bill went out of the restaurant belching*. But it corresponds to syntax noncanonically, in that (a) the manner/means constituent is expressed not as an adjunct clause or adverb but rather as the main verb; (b) *GO* in semantics is not linked to anything in syntax; (c) *way* in syntax is not linked to anything in semantics; [18] (d) the PP argument in syntax is licensed by the path argument of *GO*.

This solution contains exactly the same information as the putative lexical rule, but it uses the information differently. Instead of creating a "lexical item" of suspicious status, it effects a direct noncanonical relation between syntax and conceptual structure, supplanting the canonical linking.

(28) has two unusual properties that bear mention. First, the verb in syntax does not direct the composition of the syntactic arguments; rather, it is itself an argument of the construction, a previously unattested situation. Second, *pro's way* is not the direct object of the *verb*, since it is not a semantic argument of the verb. In fact, as noted in (27) above, the verb cannot require its own direct object if it is to participate in this construction. Rather, *pro's way* is object of the *verb phrase*, without being an argument of the verb – another previously unprecedented situation.

These conclusions may seem hard to swallow. But (a) after all, they're merely terminological; (b) a traditional derivational treatment is impossible (and I am not acquainted with any attempted derivational solutions, more than 25 years after the construction was first noticed); and (c) the lexical rule solution leads to its own uncomfortable conclusions. The interface rule solution seemed to me the most natural of the three.

If this were the only phenomenon of this sort, perhaps one wouldn't have to worry about it much. But others soon turned up, many within the Construction Grammar framework. Some have canonical syntax (29), some are noncanonical (30).

[18] Goldberg thinks it's linked to *GO*, again a noncanonical linking if there ever was one.

(29) Canonical syntax
   a. Jo slept/drank/read/programmed the afternoon/morning/evening away.[19]
      'Jo spent/wasted the afternoon sleeping' (Jackendoff 1997b)
   b. The bus rumbled around the corner.
      'The bus went around the corner in such a way as to rumble.' (Levin and Rappaport Hovav 1995; Goldberg and Jackendoff 2004)
   c. that travesty of a theory
      'that theory, which is a travesty' (Asaka 2010, Booij 2002)
(30) Noncanonical syntax
   a. Joe doesn't eat beef, let alone alligator. (Fillmore et al. 1988)
   b. The more I read, the less I understand. (Fillmore et al. 1988, Culicover and Jackendoff 2005)
   c. day by day; student after student; dollar for dollar (Jackendoff 2008)
   d. One more beer and I'm leaving. (Culicover 1972)

The ones with canonical syntax pose the same sorts of issues as the *way* construction. The ones with noncanonical syntax make the issues still clearer. In these cases, there is no evidence for an autonomous phrase structure rule in the syntax; rather there is a specialized syntactic pattern that is linked to a special semantic pattern. In other words, these constructions are at once syntactic well-formedness rules and syntax-semantics interface rules.[20]

Still, this conclusion could leave one wondering: What kind of grammar countenances rules of this sort?

9. The Parallel Architecture

Well, by now a crisis was brewing. A lot of phenomena had turned up that showed the independence of semantics from syntax (I've mentioned only a few here), and thus it became harder and harder to imagine hooking conceptual structure up to a traditional generative syntax. In particular, starting in the early 1980s, generative syntax introduced the level of Logical Form (LF), which was supposed to encode all the material in syntax relevant to semantics, in a form homomorphic to semantic structure. LF was to be derived from S(urface)-structure by "covert movement" (Chomsky 1981). However, there was never proposed a theory of semantic structure that LF was supposed to be homomorphic to. Conceptual structure didn't fill the bill, since it was not homomorphic to syntax. In fact, work on conceptual structure appeared to show that there could *not* be a level of syntax homomorphic to semantics, since these various noncanonical constructions cannot be derived by any sort of traditional movement – not to mention the much earlier finding that control in nominals cannot be derived by traditional movement and/or deletion.

Moreover, at about this time Chomsky started developing the Minimalist Program – an attempt to rebuild linguistic theory from the top down, assuming only "virtual conceptual

[19] This one arose through another epiphany: I had been teaching argument structure in an Intro class, as usual giving *sleep* as an example of a one-argument verb. And then a student asked about *sleep the night away*. Aha.
[20] In a way, this is the kind of solution that everyone has always wanted: a biunique relation between syntax and semantics. Ironically, this vision is best realized in these syntactically weird situations.

necessities." The style of analysis advocated in the MP was strikingly at odds with the phenomena of argument structure and especially the noncanonical mappings that I had been working out.

In the face of this situation, it seemed useful to step back and ask what assumptions Chomsky was making that led him to his position. This examination of the foundations of generative grammar began with *Architecture of the Language Faculty* (1997a) and continued through *Foundations of Language* (2002) to *Simpler Syntax* (2005, co-authored with Peter Culicover). Here are two of the most important assumptions; two more will be discussed in the next section.

- (Derivation) A grammar consists of a set of algorithmic principles that generate sentences step by step. This assumption made a lot of sense in 1957, but by 1990 there were many alternative frameworks on the market that characterized grammars as sets of simultaneously applied constraints.
- (Syntactocentrism) Syntax is the only generative component of the grammar; phonological and semantic structures are derived from syntax. This assumption, explicit in *Aspects* and never ever defended, quickly became an unnoticed dogma of generative theory.

These assumptions are so ingrained in the mainstream mindset that it is nearly impossible to dislodge them. In particular, syntactocentrism requires every systematic semantic distinction to be derived from a syntactic distinction. It thereby forces syntactic theory to countenance a massive covert complexity, which practitioners become accustomed to touting as evidence for the abstract complexity of the language faculty.

It is only through the development of a robust theory of meaning such as Conceptual Semantics that it is possible to see the alternative.[21] Granting semantics its own generative capacity, with a more flexible interface to syntax, absolves syntax of full responsibility for meaning. Simpler Syntax sets a different goal for simplicity than does mainstream syntax: syntax should consist of the minimum structure necessary to map between semantics and phonology. As shown in *Simpler Syntax*, this permits a much simpler account of virtually every major syntactic phenomenon, as well as of many phenomena not recognized in the mainstream syntactic canon. In particular, Simpler Syntax has no movement and no empty categories other than perhaps A'-traces. It has relatively flat, multiply branching trees rather than MP's deep binary branching trees. Given these factors, it is more "surfacey" than most frameworks of grammar.[22]

[21] For instance, a prominent mainstream syntactician recently said to me, "I'm not sure I believe in conceptual structure. Why should language have two structures, syntax and conceptual structure, that are so much alike?" This question comes from a lifetime of syntactic theory in which syntax is pushed ever closer to (presumed) semantics by the need to derive semantics from it. It also comes from the mainstream's disregard for lexical semantics, a predilection shared by and reinforced by formal semantics.

[22] Though it does have a grammatical function tier, a very limited version of LFG's f-structure (Bresnan 1982, 2001). Culicover and I found ourselves forced to this conclusion, noting however that practically every major theory of syntax has a counterpart, such as abstract Case in mainstream syntax, not to mention the entire theory of Relational Grammar.

In arriving at the position that conceptual structure has its own autonomous structure, linked to but not entirely determined by syntax, I was heartened by what I saw as parallel developments in phonology in the 1970s and 1980s (as mentioned already in section 4). Phonotactic rules play much the same role in phonology as phrase structure rules do in syntax, establishing a set of permissible phonological forms. Prosodic phonology divides phonological structure into syllables, feet, and larger constituents such as intonational phrases, which do not correlate directly with syntactic constituents (Selkirk 1983). It proves pointless to derive these phonological constituents by the "readjustment rules" of SPE. Rather, it makes more sense to posit a flexible interface that links phonological and syntactic constituency.[23] Autosegmental phonology (Goldsmith 1976) divides phonological structure into independent tiers (segmental/syllabic tier, metrical/stress grid, intonation contour, tone tier), each with its own characteristic structure. None of them is derived from the others; rather, each of them is linked to the others by a flexible interface. For example, word stress is the result of a compromise between an ideal metrical grid (optimal autonomous metrical form) and syllable weight (through the segmental-metrical interface).

Looking at these developments through the lens of the reconceptualization of the relation between syntax and semantics, it became clear that the very same relation holds between syntax and phonology (although hardly anyone in phonology was drawing the same heretical conclusion I did). Phonology has its own generative capacity – in fact multiple generative capacities – and it is linked to syntax by a many-to-many interface. This position may be controversial in semantics, given the overall level of controversy about meaning. But in the domain of phonology, where we have the best analysis of any component of language, the evidence is altogether compelling (and was immediately embraced by phonologists).

Another confirming consideration was the relation of conceptual structure to spatial structure. As described in section 5, these two structures are built out of altogether different sorts of units, but they are coordinated by an interface that links some aspects of conceptual structure (those that encode spatial concepts) to some aspects of spatial structure (those that pick out objects, their parts, and their spatial relations, but not those that delineate exact shape).

At the other end of language, phonetic structure must be related to both auditory processing and motor control of the vocal tract. Each of these involves aspects not present in the others. Neither auditory nor motor representations include such phonological notions as word boundaries and discrete stress levels. In the other directions, the auditory signal codes the speaker's voice quality, tone of voice, and speed of articulation, none of which are part of phonetics. Similarly, the motor signal can be modulated by holding an object in one's teeth. Thus each of these representations has properties that cannot be derived from the others.

[23] One of the howlers in *Aspects* was Chomsky's assertion that the normal intonation of *This is the cat that ate the rat that ate the cheese* is a performance error, and that the "correct" intonation would follow the right-embedded syntax. Prosodically, the flat three-part intonation of this sentence is perfect, while a prosody that mimicked the syntax would be highly unnatural. This speaks for the semi-independence of prosodic structure from syntactic structure.

The overall conclusion is that the connection between different components of language is qualitatively of the same sort as the connections between language and the rest of the mind, both at the conceptual and at the sensory-motor end.

The upshot is an utter rejection of syntactocentrism . It is also a rejection of derivation as the sole source of linguistic structure. Interfaces do not *derive* one kind of structure from another: they simply state relations that may or must obtain between various parts of two (or more) independent structures.

So here is the canonical picture of the Parallel Architecture. It can be seen as a close-up of part of the overall architecture of mind in section 5.

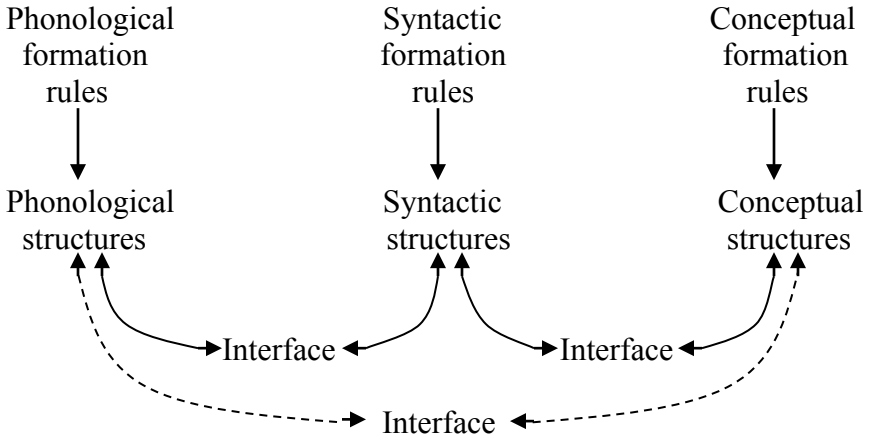

Figure 2. The Parallel Architecture

The structure of a sentence consists of well-formed structures in all three domains, linked in well-formed fashion by interface rules.

We've already spoken of the phonology-syntax interface and the syntax-semantics interface. Figure 2 includes also a tentative phonology-semantics interface. This was originally included in order to represent a direct connection between stress and intonation in phonology on one hand and information structure (topic and focus) on the other. In English, these can be independent of syntactic structure:

(31) a. The GIRL chased the boy.
b. The girl CHASED the boy.
c. The girl chased the BOY.

A syntactocentric theory has to code topic and focus in syntax, perhaps through some sort of dummy feature, even though there is no necessary syntactic reflex.[24] In a Parallel Architecture, the relation between prosody and information structure can bypass syntax altogether, while leaving open the additional possibility of constructions such as clefts that do mark information

[24] And this is how I did it in *Semantic Interpretation in Generative Grammar* (1972), which presumed a syntactocentric architecture.

structure syntactically. As it turned out, the phonology-semantics interface proves useful in other respects as well (see section 11).

An architecture based on independent interacting components was not without precedent in linguistic theory. As just mentioned, the internal organization of phonology has this character. The architecture of LFG (Bresnsn 1982, 2001) involves two independent and interacting tiers in syntax: c-structure and f-structure – and more tiers in recent versions. Sadock's (1991) Autolexical Syntax proposes independent phrasal and morphological tiers in syntax, with interface rules between them – in some versions further subdivided into subtiers; Van Valin's (1997) Role and Reference Grammar distinguishes independent and interacting levels of phrasal syntax, morphosyntax, semantic argument structure and information structure. In addition, as mentioned in section 4, the musical grammar that Lerdahl and I developed has this overall form. So although the Parallel Architecture represented a major departure from mainstream generative grammar, there were quite a few other trains departing on the same track.[25]

10. The nature of lexical items

At the same time, there was another crisis in the works: the treatment of the lexicon. As mentioned in section 8, several issues arose in how to think of "lexical rules" and how to think of the status of constructions like the *way* construction.

But a larger set of issues arose rather fortuitously. My then 10-year-old daughter Beth was addicted to the television show *Wheel of Fortune*, in which guests win prizes for guessing mystery words and phrases. I was not quite sure why, but it struck me as interesting to get an idea of the range of these items, so I asked her to collect the solutions as she watched. After a few months, she came up with a list of about 600 items. And then I was unexpectedly invited to a conference on idioms in Tilburg, and only then did I realize just why this was important: speakers of American English *know* all these idioms (*a breath of fresh air*, *food for thought*), names (*Walla Walla, Washington; American Heart Association*), clichés (*a day that will live in infamy*, *high as a kite*), quotations (*and the rockets' red glare*, *may the Force be with you*), and so on. That is, they must be stored in memory – and there are thousands upon thousands of them.

How should they be stored? One possibility, which turns up in Chomsky's work from time to time, is that idioms are X^0 categories with internal syntactic structure. This is improbable for a number of reasons. For instance, if *a breath of fresh air* is a noun, it should be able to occur with noun modifiers, e.g. **the lovely a breath of fresh air*. And it's a coincidence that inside of nouns there can be normal relative clauses, such as *a day that will live in infamy*. Furthermore, a discontinuous idiom such as *give NP a hard time* either has to allow a syntactic argument inside a V^0 or else extrapose the verb-internal NP to – surprise! – the very position where such an argument would normally appear in phrasal syntax. I take it that this position is hopeless (though I recognize that some people will not find it so).

[25] Though none of the others made reference to the way the language capacity is embedded in the rest of the mind, for me a central point.

Another sort of treatment of these items is to delegate them to a separate storage from words, sometimes called a "constructicon" or the like. But notice: like ordinary words, they have phonology and semantics. They differ from words only in that they are syntactically composite. Thus there is no need to "put them in a different place"; they are intrinsically distinguished by their structure. We can call them phrasal lexical items.

However, the terminology here points up an inherent problem in the conception of the lexicon handed down to standard generative grammar from traditional grammar:

- The grammar captures all the regularities in the language; the lexicon is a list of exceptions.
- The lexicon consists of words (or morphemes).

A breath of fresh air violates both these assumptions. It is not completely exceptional, since it is made up of existing words with standard syntax. Moreover, it is an entire NP, not just a word. If we want to include it in the lexicon, then, we need to alter our conception of lexical items.

This problem was addressed by DiSciullo and Williams (1987), who made a distinction between grammatical words (X^0 categories in syntax, phonological words in phonology) and what they called "listemes" – the stored items, whatever their structure. Their contention is that the theory of listemes is "of no interest to the grammarian." Wrong! The grammarian (at least in generative grammar) is trying to characterize "knowledge of language," which implies things stored in memory. If we want to take this goal seriously, we therefore have to ask what a speaker stores, whether words or phrases.[26]

To make it easier to keep the distinction in mind, I'll use the term *lexical* to mean 'pertaining to the lexicon,' and *lexical item* to mean 'something stored in the lexicon.' I'll distinguish the structure of words (the domain of morphology) from that of phrases (the domain of phrasal syntax); and in parallel to the term *phrasal* ('pertaining to phrases'), I'll coin the term *wordal* ('pertaining to words'). The usual use of the term *lexical* conflates 'stored in the lexicon' and 'pertaining to words'; *wordal* replaces the latter use. We can then say, for instance, that *dog* is both lexical and wordal; *a breath of fresh air* is lexical but phrasal; *a whiff of rotten potatoes* is phrasal but not lexical.

To make this observation more systematic, we can ask what a word is. In any mentalistic theory, it is a long-term memory linkage of a word-sized piece of phonology, a set of syntactic features (including that it belongs to a zero-level category), and a piece of meaning. In classical generative grammar, words are inserted into syntax by "lexical insertion" or, in the MP, by Merge. Their phonology and semantics are carried around passively through the syntactic derivation and then "read off" when the derivation is "passed" to the appropriate component. The Parallel Architecture has a different story: phonological and semantic features have no place in the syntactic component. Rather, a word can be thought of as a little interface rule, licensing a link between pieces of structure in the three components. If you want to think in terms of "lexical insertion," the three pieces of the word are inserted in parallel, adding their links to the

[26] And this came to be a rather burning issue in psycholinguistics as well, thanks to the "words vs. rules" controversy (e.g. Rumelhart and McClelland 1986, Pinker and Prince 1988, Pinker 1999, Nooteboom, Weerman, and Wijnen 2002).

relation among the three structures. If you want to think in terms of licensing constraints, a word licenses pieces of structure in the three components, plus well-formed links among them.

This leads to an important difference from traditional syntactic notation. Here's a phrase in traditional notation:

(32)

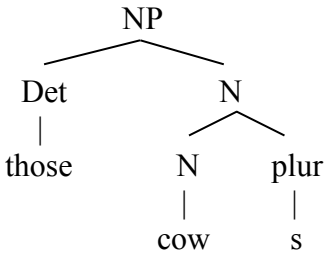

But in the Parallel Architecture, it's important to keep phonological, syntactic, and semantic structures distinct. The fact that the words are *those* and *cow* rather than something syntactically identical such as *these* and *dog* is not a matter for syntax – all the relevant distinctions are in phonology and semantics. So a better notation rigorously segregates the three components and notates their linkings by subscripts:

(33) Phonological structure Syntactic structure Conceptual structure

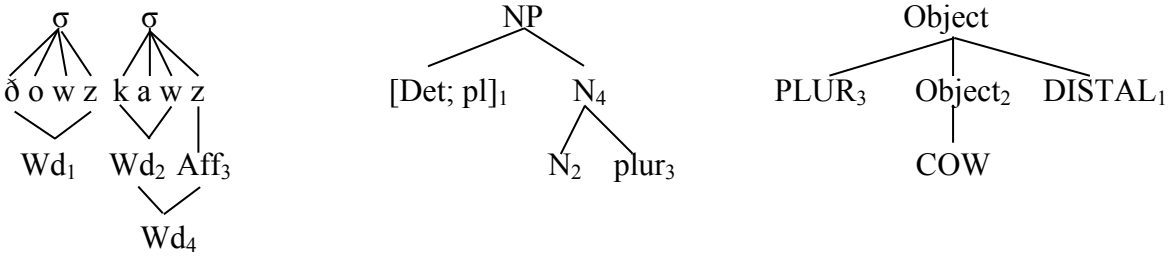

If we want to use the traditional notation, we have to understand it as an abbreviation for the tripartite notation in (33).

11. The inhabitants of the lexicon

An interesting advantage of the Parallel Architecture treatment was pointed out by my then-student Urpo Nikanne: there are words of English that have phonology and semantics but no apparent syntax, for instance those in (34).

(34) hello, goodbye, ouch, yes, oops, gosh, dammit, upsy-daisy, allakazam, feh, uh-oh, wow, gee whiz, oboy, yuck, ugh, phooey, gadzooks, fiddlesticks, oo-la-la, yikes

Unlike ordinary words, these can function as complete utterances. Moreover, they can't be combined with other words into sentences, except by parataxis (35a) and in quotative and metalinguistic contexts where anything can be inserted, even a phrase of another language (35b,c).

(35) a. Hello, Bill.
Hello, I must be going.
b. "Hello," she said. (cf. *"La plume de ma tante," she said.*)
c. the lexical item 'hello' (cf. *the Dutch word 'lekker'*)

In a syntactocentric grammar, these have to have some syntactic structure or other into which they are inserted: phonology and semantics can't be linked if there's no syntax to generate them. In a Parallel Architecture, they are simply linkages of phonological and conceptual structure, lacking syntactic features altogether.

(36) Semantics: $[_{\text{Event}}$ I GREET YOU$]_5$
Phonology: $[_{\text{Wd}}$ hello$]_5$

We can go even farther: there are even words (or whatever you want to call them) that are just stored pieces of phonology without either syntax or meaning. These are nonsense refrains used to fill up metrical structure in songs and nursery rhymes:

(37) fiddle-de-dee, hey-diddle-diddle, hickory-dickory-dock, eenie-meenie-minie-moe, she-bop-she-bop, rickety-tickety-tin, ink-a-dink-a-doo, e-i-e-i-o, brrr-raka-taka

One might be tempted to say these aren't language. But what are they then? They have phonological structure and stress, and they occur intermixed with text. If you know the songs, you know these items. Rather than invent a new "nonsense-wordicon," the simplest thing to do is to put them in the lexicon along with ordinary words. The fact that they lack any syntax or semantics guarantees that they won't appear in ordinary sentences. That is, like idioms, their intrinsic properties account for their distribution, and there is no need to further distinguish them in the theory.

A different combination of properties appears in a limited number of words that have syntax and phonology but no meaning. They function only as "grammatical glue." (In some versions of syntactocentric grammar, their status is recognized by inserting them late in the syntactic derivation – especially in the case of *do*-support.) (39) is a lexical entry for the *of* in (38c).

(38) a. <u>It</u>'s hot in here. <u>It</u> seems as though <u>it</u>'s going to rain.
b. <u>Do</u> you want a drink? We <u>do</u> not allow dogs in here.
c. a picture <u>of</u> Bill
d. I think <u>that</u> it's going to rain.

(39) Syntax: P_6
Phonology: of_6

If one believes in PRO (or *pro*), these are items that have syntax and semantics and a link between them, but that lack phonology. Simpler Syntax eschews PRO and *pro* but has other items with syntax and semantics but no phonology – see below.

The other logical possibilities are items with syntax alone and items with semantics alone. We will also see such cases in a moment.

(34)-(39) illustrate word-sized lexical items that have structure in only one or two of the parallel components. Earlier in this section we saw phrase-sized lexical items such as *a breath of fresh air*, which link phonology, syntax, and semantics. So now consider idioms with noncanonical structure such as those in (40).

(40) day in day out
by and large
if need be
for the time being
through and through[27]
How about XP?

Although perhaps the individual words belong to syntactic categories, there is no evidence that they are connected syntactically – there is no syntactic principle that licenses their connection. So one might think of these as lexical items with phonology and semantics but "defective" syntax.

Consider next the variables involved in specifying argument structure. If a verb has a semantic argument but does not require it to be expressed in syntax (e.g. *eat*, *read*), the variable is a piece of semantics unconnected to syntax or phonology. If the semantic argument must be expressed in syntax, i.e. it is obligatorily subcategorized (e.g. *devour*), then the semantic and syntactic variables are present and linked, but no phonology is specified. The argument of a discontinuous idiom (e.g. *take NP to task*) again specifies linked syntactic and semantic variables, but no phonology.

Returning to the constructional idioms discussed in section 8, these now can be treated as lexical items with perfectly well-formed syntax and semantics, but an unusual configuration of linking between syntactic and semantic variables. In the *way* construction, for instance, only one syntactic piece, the noun *way*, is linked to phonology – and this piece is not linked to semantics! So the oddity of this construction is reflected in the oddity of its formalization, without violating the overall guidelines for the format of lexical items. (41) is almost identical to (28), with one further refinement: in keeping with the Parallel Architecture, *way* belongs in phonological structure, linked to a noun in syntax (the arrows could be replaced with subscripts).

[27] According to Booij 2010, 59-60, this construction is slightly more productive in Dutch, where it is an intensifier attached to adjectives, e.g. *door- en door-nat*, *in- en in-triest*, 'through and through wet,' 'in and in sad.'

(41) Semantics: [$_{Event}$ GO (X, Path); Manner/Means: [$_{Event}$ F (X)]]

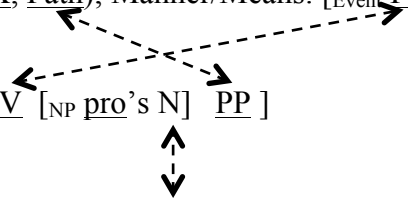

Syntax: [$_{VP}$ V [$_{NP}$ pro's N] PP]

Phonology: way

More generally, the standard semantics-syntax linking rules (e.g. the default linking of Agent to subject position) are schemas that link variables in conceptual structure to variables in syntax. In other words, they too can be stated in the same format and therefore can qualify as lexical items as well.

All these cases involve linkings among two of the three components of the grammar. Could there be lexical items that consist only of syntactic structure? Yes. Phrase structure rules are typically stated in derivational form, as in (42a), the rule for the English transitive VP. A notational variant of (42a) specifies the form of the VP as a template or schema that can be used to license syntactic structures, as in (42b). Janet Fodor has called such schemas "treelets"; Tomasello 2003 calls them "slot and frame schemas."

(42) a. VP → V – (NP) – …
b. [$_{VP}$ V (NP) …]

(42b) states a condition to which English VPs have to conform, whatever they mean. It is thus a rule of autonomous syntax. And yet it has the same format as all the other lexical items we've encountered.

How can (42b) belong in the lexicon? It's a rule of grammar! Well, we've seen lexical items consisting only of phonology, items that link syntax and phonology alone, items that link semantics and phonology alone, and items with variables in one or more domains. (42b) just fills another cell in this multidimensional space. In fact, the theory would be *less* general if it excluded phrase structure rules from the lexicon.[28] On this view, then, the formation rules in Fig. 2 as well as the interface rules are part of the lexicon.

The outcome is a view of language quite different from what we're used to – and it actually makes some things work better. In particular, we can easily incorporate all these intermediate cases such as idioms and meaningful constructions, whereas standard approaches are more or less at a loss, and haven't come up with a plausible account of them. Instead of a dichotomy between grammar and lexicon, there is a three-dimensional continuum:

- Is the lexical item wordal (*dog*) or phrasal (*a breath of fresh air*)?

[28] This approach only works, though, in a syntactic theory that doesn't have movement. Movement rules can't be stated in this format – especially rules that apply in some algorithmic order.

- What combination of phonological, syntactic, and conceptual structure does the item specify? *Dog* and *a breath of fresh air* stipulate all three; the items in (34)-(40) and (42b) stipulate one or two of the structures; (41) stipulates a mixture of linkages.
- Does the lexical item contain variables? *Dog* does not; some of the material in discontinuous idioms and constructions like (41) consists of variables; a phrase structure rule like (42b) consists entirely of variables.

In each case, a lexical item is simply a piece of linguistic structure. The traditional lexicon represents only one corner in the continuum. Given its radical expansion over the traditional lexicon, maybe we should have a different name for this version (my students liked Lexicon 2.0). I should also mention that similar arguments have been made in the context of Cognitive Grammar, Construction Grammar, and HPSG.

12. Further frontiers

From the point of view of standard generative grammar, this conception of language looks kind of weird, and I have to admit that it took me several years to feel comfortable with it. But it is now beginning to pay off, especially in its ability to treat words and rules as forming a continuum.

Here's one example. Consider differences in register. It is easy to add a feature to words to mark whether they are formal, casual, or vulgar. But now, since syntactic constructions are also lexical items, they can be marked for register in the very same way. For instance, the syntax of a pied-piped wh-construction such as *To whom are you speaking* involves a little treelet in the lexicon of the form (43), and this can be marked for formal register just like a word such as *assuredly*. There is no longer a need to condition a movement rule on register (a la Labov 1969) – an entirely different mechanism than for words.

(43) Syntax: $[_S\ [_{PP}\ P\ [_{NP}$ wh-]] …]
Register: formal

The same goes for dialect differences – dialect-specific words, phonetic variants, and syntactic structures can all be labeled in a uniform way – there aren't different kinds of markings for words, phonology, and syntactic derivation.

For another example: I've talked so far about word-sized and phrase-sized lexical items. But the same treatment can be applied to items smaller than a word, for instance regular morphology. The English regular plural affix can be encoded like this (and this was implied in example (33)):

(44) Semantics: $[PLUR_2\ (\underline{X}_1)]$
Syntax: $[_N\ \underline{N}_1 - plur_2\ ]$
Phonology: $[_{Wd}\ \underline{Y}_1 - z_2]$

More generally, the approach of Construction Morphology (Booij 2010) lends itself to being incorporated into this framework.

For a different sort of example: standard generative grammar is notorious for making a strong distinction between theories of competence – the grammar – and theories of performance – language comprehension and production, and saying that the connection between the two is mysterious and probably not the business of linguistic theory (though this is beginning to change among some proponents). In contrast, the Parallel Architecture translates directly into a theory of language processing (Jackendoff 2002, 2007b). The basic idea comes from the account of words as interface rules. For instance, when you hear a string of sounds, this calls up words that sound like this string, setting up a phonological structure. But since the word is an interface rule, it also brings along its syntax and meaning, which help construct syntactic structures and conceptual structures. Similarly, in production, a concept you want to express invokes one or more words through their semantics, but then they drag along their syntax and phonology, which help construct a pronunciation for the thought to be expressed. Within this approach, some of my students and colleagues have been pursuing research on the processing of some of these constructions where the semantics doesn't correlate in the normal way with the syntax, and we've gotten some interesting results (e.g. Piñango, Zurif, and Jackendoff 1999; Kuperberg et al. 2010; Wittenberg et al. in press).

The Parallel Architecture also lets us ask the question: What if there were no syntax, and you had to map directly between phonology and meaning? In the standard approach, this would be impossible, because phonology and meaning can't exist without syntax. But if they're independent generative systems, there can be interface principles that bypass syntax altogether. In work I've been pursuing with Eva Wittenberg (Jackendoff and Wittenberg in press), we've found that this begins to give us a way to describe how less complex linguistic systems work, such as early child language, pidgins, and emerging languages like Al-Sayyid Bedouin Sign Language. These systems *do* manage to convey meaning, albeit perhaps with less sophistication that full languages. So the Parallel Architecture opens up many new connections to phenomena that the standard approach can't touch.

To sum up, I started by developing a theory of conceptual structure that was robust enough to stand on its own. This eventually forced me diverge from standard generative grammar, in accepting the independence of semantics from syntax. In turn, this led to the following conclusions:

- Semantics is in a many-to-many relation with syntax, not one-to-one.
- Since syntax is no longer responsible for everything in semantics, it becomes much simpler.
- Semantics can be seen as one of the central mental representations involved in conceptualizing the world and in reasoning.
- Conceptual structure appears in nonlinguistic organisms such as babies and apes, so it is not dependent on language.
- It interfaces not only with language but with visual and other kinds of perception.
- Phonology, syntax, and semantics are independent generative systems linked by interface rules.

- Words are interface rules.
- There is no strict division between the lexicon and the grammar; rather words and standard rules are at opposite corners of a multidimensional continuum that includes all sorts of mixed items such as idioms and meaningful constructions.

This represents a major reconceptualization of the nature of language and the language faculty that permits the integration of the disparate components of language, their integration with other faculties of mind, and their participation in language processing.

Some critics of the Parallel Architecture (e.g. Phillips and Lau 2004, Marantz 2005) have faulted it for not being sufficiently constrained. The system against which it is being implicitly compared is the Minimalist Program, which is built on the premise that above all an account of the language faculty must make use of the least possible resources. The difficulty with this view is that it hobbles the investigator from the start. This is taken to be a virtue of the methodology. Yet the result is a theory that excludes many robust syntactic phenomena such as idioms and meaningful constructions, that has no theory of phonology, semantics, or the lexicon, that makes no contact with language processing, and that makes no contact with non-linguistically supported thought. It also bears mention that the syntactic component of the Parallel Architecture is in fact highly constrained, though in a different way: it has no movement and almost no empty nodes, precisely the factors in classical generative grammar that engender the greatest need for constraint.

I maintain that although general scientific grounds motivate an attempt to constrain the theory appropriately, two larger goals are of paramount importance: (a) describing all the facts, and (b) connecting with related scientific inquiries. As I said in *Foundations of Language*, one of the reasons linguistics has become estranged from the other cognitive sciences is that it has systematically ignored their issues. I hope that the increased interest among linguists in experimental research is going to help change this. On the other hand, much of the experimentation is conducted in a somewhat atheoretical context. So I want to end by emphasizing the importance of theory to the scientific enterprise. I hope the Parallel Architecture will be a useful vehicle for framing research to come.

References

Asaka, T. 2002. A lexical licensing analysis of the adjectival noun construction. *English Linguistics* 19, 113-141

Baker, Mark. 1988. *Incorporation: A Theory of Grammatical Function Changing*. Chicago: University of Chicago Press.

Barkow, Jerome, Leda Cosmides, and John Tooby (eds.). 1992. *The Adapted Mind*. New York: Oxford University Press.

Bernstein, Leonard. 1976. *The Unanswered Question*. Cambridge, MA: Harvard University Press.

Berwick, Robert, and Noam Chomsky. 2011. The Biolinguistic Program. The current state of its development. In Di Sciullo and Boeckx 2011, 19-41.

Bloom, Paul, Mary Peterson, Lynn, Nadel, and Merrill Garrett (eds.). 1996. *Language and Space*. Cambridge, MA: MIT Press

Booij, Geert. 2002. Constructional idioms, morphology and the Dutch lexicon. *Journal of Germanic Linguistics* 14, 301-329.

Booij, Geert. 2010. *Construction Morphology*. Oxford: Oxford University Press.

Bresnan, Joan (ed.). 1982. *The Mental Representation of Grammatical Relations.* Cambridge, MA: MIT Press
Bresnan, Joan. 2001. *Lexical Functional Syntax*. Oxford: Blackwell.
Byrne, Richard, and Andrew Whiten (eds.). 1988. *Machiavellian Intelligence*. Oxford: Clarendon Press.
Carey, Susan. 2009. *The Origin of Concepts*. Oxford University Press.
Cheney, Dorothy, and Robert Seyfarth. 1990. *How Monkeys See the World.* Chicago: University of Chicago Press.
Cheney, Dorothy, and Robert Seyfarth. 2007. *Baboon Metaphysics*. Chicago: University of Chicago Press.
Chomsky, Noam. 1965. *Aspects of the Theory of Syntax*. Cambridge, MA: MIT Press.
Chomsky, Noam. 1968. *Language and Mind*. New York: Harcourt, Brace & World.
Chomsky, Noam. 1981. *Lectures on Government and Binding*. Dordrecht: Foris.

Chomsky, Noam, and Morris Halle. 1968. *The Sound Pattern of English*. New York: Harper & Row.
Csibra, G., S. Biro, O. Koos, and G. Gergely. 2003. One-year-old infants use teleological representations of actions productively. *Cognitive Science* 27, 111-133.
Culicover, Peter W. 1972. OM-sentences: On the derivation of sentences with systematically unspecifiable interpretations. *Foundations of Language* 8, 199-236.
Culicover, Peter W., and Ray Jackendoff. 2005. *Simpler Syntax*. Oxford: Oxford University Press.
Davidson, Donald. 1967. The logical form of action sentences. In N. Rescher (ed.), *The Logic of Decision and Action*. Pittsburgh: University of Pittsburgh Press.
de Waal, Frans. 1996. *Good Natured*. Cambridge, MA: Harvard University Press.
DiSciullo, Anna Maria, and Edwin Williams. 1987. *On the Definition of Word*. Cambridge, MA: MIT Press.
Fillmore, Charles, Paul Kay, and Mary Catherine O'Connor. 1988. Regularity and idiomaticity in grammatical constructions. *Language* 64, 501-53.
Fodor, Jerry. 1970. Three reasons for not deriving 'kill' from 'cause to die.' *Linguistic Inquiry* 1, 429-438.
Fodor, Jerry. 1975. *The Language of Thought*. Cambridge, MA: Harvard University Press.
Goldberg, Adele. 1995. *Constructions.* Chicago: University of Chicago Press.
Goldberg, Adele, and Ray Jackendoff. 2004. The English resultative as a family of constructions. *Language* 80, 532-568.
Goldsmith, John. 1979. *Autosegmental Phonology*. New York: Garland Press.
Golinkoff, Roberta, and Kathy Hirsh-Pasek. [knowledge of argument structure]
Gruber, Jeffrey. 1965/1970. *Studies in Lexical Relations*. Ph.D. dissertation, MIT; reprinted in Gruber, *Lexical Structures in Syntax and Semantics*, Amsterdam, North-Holland.
Hale, Kenneth, and Samuel Jay Keyser. 1993. On argument structure and the lexical expression of syntactic relations. In K. Hale and S. J. Keyser (eds.), *The View from Building 20*, 53-109. Cambridge, MA: MIT Press.
Hankamer, Jorge, and Ivan Sag. 1976. Deep and surface anaphora. *Linguistic Inquiry* 7, 391-428.
Hauser, Marc. 2006. *Moral Minds: How Nature Designed our Sense of Right and Wrong*. New York: HarperCollins/Ecco Press.
Herskovits, Annette. 1986. *Language and Spatial Cognition*. Cambridge, Cambridge University Press.
Hirsh-Pasek, Kathryn, and Roberta Golinkoff. 1996. *The Origins of Grammar*. MIT Press.
Jackendoff, Ray. 1972. *Semantic Interpretation in Generative Grammar*. Cambridge, MA: MIT Press.
Jackendoff, Ray. 1974. A deep structure projection rule. *Linguistic Inquiry* 5, 481-506
Jackendoff, Ray. 1983. *Semantics and Cognition*. Cambridge, MA: MIT Press.
Jackendoff, Ray. 1987. *Consciousness and the Computational Mind*. Cambridge, MA: MIT Press/Bradford Books.
Jackendoff, Ray. 1990. *Semantic Structures*. Cambridge, MA: MIT Press.

Jackendoff, Ray. 1992. *Languages of the Mind*. Cambridge, MA: MIT Press.
Jackendoff, Ray. 1994. *Patterns in the Mind*. New York: Basic Books.
Jackendoff, Ray. 1997a. *The Architecture of the Language Faculty*. Cambridge, MA: MIT Press.
Jackendoff, Ray. 1997b. Twistin' the night away. *Language* 73, 534-559.
Jackendoff, Ray. 2002. *Foundations of Language*. Oxford: Oxford University Press.
Jackendoff, Ray. 2007a. *Language, Consciousness, Culture*. Cambridge, MA: MIT Press.
Jackendoff, Ray. 2007b. A Parallel Architecture Perspective on Language Processing, *Brain Research* 1146, 2-22.
Jackendoff, Ray. 2008. *Construction after construction* and its theoretical challenges. *Language* 84, 8-28.
Jackendoff, Ray. 2011. What is the human language faculty? Two views. *Language* 87, 586-624.
Jackendoff, Ray, and Eva Wittenberg. in press. What you can say without syntax: A hierarchy of grammatical complexity. In Fritz Newmeyer and Lauren Preston (eds.) …. Oxford: Oxford University Press.
Johnson, Mark. 1987. *The Body in the Mind*. University of Chicago Press.
Katz, Jerrold, and Jerry Fodor. 1963. The structure of a semantic theory. *Language* 39, 170-210.
Katz, Nancy, Erica Baker, and John Macnamara. 1974. What's in a name? A study of how children learn common and proper names. *Child Development* 45, 469-473.
Köhler, Wolfgang. 1927. *The Mentality of Apes*. New York: Vintage Books.
Kosslyn, Stephen. 1980. *Image and Mind*. Cambridge, MA: Harvard University Press.
Kuperberg, Gina, Arim Choi, Neil Cohn, Martin Paczynski, and Ray Jackendoff. 2010. Electrophysiological correlates of Complement Coercion. *Journal of Cognitive Neuroscience* 22(12): 2685-2701.
Labov, William. 1969. Contraction, deletion, and inherent variability of the English copula. *Language* 45, 715-762.
Lackner, James. 1981. Some contributions of touch, pressure, and kinesthesis to human spatial orientation and oculomotor control. *Acta Astronautica* 8, 825-830.
Lackner, James, and Paul Dizio. 2000. Aspects of body self-calibration. *Trends in Cognitive Sciences* 4, 279-288.
Lakoff, George. 1965/1971. *On Syntactic Irregularity*. New York: Holt, Rinhart & Winston
Lakoff, George. 1970. On generative semantics. In Steinberg and Jakobovits 1970
Lakoff, George, and Mark Johnson. 1980. *Metaphors We Live By*.
Landau, Barbara, and Lila Gleitman. 1985. *Language and Experience: Evidence from the Blind Child*. Cambridge, MA: Harvard University Press.
Landau, Barbara, and James Hoffman. 2012. *Spatial Representation*. Oxford University Press.
Landau, Barbara, and Ray Jackendoff. 1993. "What" and "where" in spatial language and spatial cognition. *Behavioral and Brain Sciences* 16, 217-238. (Shorter version in Jackendoff 1992)
Landau, Barbara, Linda Smith, and S. Jones. 1988. The importance of shape in early lexical learning. *Cognitive Development* 3, 299-321.
Langacker, Ronald. 1987. *Foundations of Cognitive Grammar*. Stanford: Stanford University Press.
Lerdahl, Fred, and Ray Jackendoff. 1982. *A Generative Theory of Tonal Music*. Cambridge, MA: MIT Press.
Levelt, W. J. M. 1989. *Speaking*. Cambridge, MA: MIT Press.
Levin, Beth, and Malka Rappoport Hovav. 1995. *Unaccusativity: At the Syntax-Lexical Semantics Interface*. Cambridge, MA: MIT Press.
Liberman, Mark, and Alan Prince. 1977. On stress and linguistic rhythm. *Linguistic Inquiry* 8, 249-336.
Luo, Yuyan, and Renée Baillargeon. 2005. Can a self-propelled box have a goal? Psychological reasoning in 5-month-old infants. *Psychological Science* 16, 601-608.
Macnamara, John. 1978. How do we talk about what we see? Ms., McGill University.
Marantz, Alec. 2005. Generative linguistics within the cognitive neuroscience of language. *The Linguistic Review* 22.429–46.

Marr, David. 1982. *Vision.* San Francisco: Freeman.
Marr, David, and H. K. Nishihara. 1978. Visual information processing: Artificial intelligence and the sensorium of sight. *Technology Review* 81 (October), 28-49.
McCawley, James. 1968. Lexical insertion in a transformational grammar without deep structure. CLS 4.
McCawley, James. 1978. Conversational implicature and the lexicon. In P. Cole (ed.), *Syntax and Semantics 9*, 245-259. New York: Academic Press.
Miller, George, and Philip Johnson-Laird. 1976. *Language and Perception*. Cambridge, MA: Harvard University Press.
Nootebohm, Sieb, F. Weerman, and F. Wijnen (eds.). 2002. *Storage and Computation in the Language Faculty.* Dordrecht: Kluwer.
Phillips, Colin, and Ellen Lau. 2004. Foundational issues (Review article on Jackendoff 2002). *Journal of Linguistics* 40.571–91.
Piñango, Maria, Edgar Zurif, and Ray Jackendoff. 1999. Real-Time Processing Implications of Enriched Composition at the Syntax-Semantics Interface. *Journal of Psycholinguistic Research*, 28, 395-414.
Pinker, Steven. 1989. *Learnability and Cognition*. Cambridge, MA: MIT Press.
Pinker, Steven. 1999. *Words and Rules*. New York: Basic Books.
Pinker, Steven, and Alan Prince. 1988. On language and connectionism. *Cognition* 26, 195-267.
Povinelli, Daniel. 2000. *Folk Physics for Apes*. Oxford: Oxford University Press.
Rosenbaum, Peter. 1967. *The Grammar of English Predicate Complement Constructions.* Cambridge, MA: MIT Press.
Rumelhart, David, and James McClelland. 1986. On learning the past tense of English verbs. In J. McClelland, D. Rumelhart, and the PDP Research Group, *Parallel Distributed Processing, Vol. ii*, 216-271. Cambridge, MA: MIT Press.
Sadock, Jerrold. 1991. *Autolexical Syntax*. Chicago: University of Chicago Press.
Schenker, Heinrich. 1935. *Der freie Satz*. Vienna: Universal Edition. English translation by Ernst Oster: New York, Longmans, 1972.
Selkirk, Elizabeth. 1984. *Phonology and Syntax*. Cambridge, MA: MIT Press.
Song, Hyun-joo, Renée Baillargeon, and Cynthia Fisher. 2005. Can infants attribute to an agent a disposition to perform a particular action? *Cognition* 98 2005, B45-B55.
Talmy, Leonard. 1978. The relation of grammar to cognition: A synopsis. In D. Waltz (ed.), *Theoretical Issues in Natural Language Processing 2*, 14-24. New York: Association for Computing Machinery.
Talmy, Leonard. 1983. How language structures space. In H. Pick and L. Acredolo (eds.), *Spatial Orientation: Theory, Research, and Application*. New York: Plenum.
Tomasello, Michael (ed.) 2000. Primate cognition. Special issue of *Cognitive Science* 24.
Tomasello, Michael. 2003. *Constructing a Language.* Cambridge, MA: Harvard University Press.
Ungerleider, L. G., and M. Mishkin. 1982. Two cortical visual systems. In D. J. Ingle, M. A. Goodale, and R. J. W. Mansfield (eds.), *Analysis of Visual Behavior*, 549-586. MIT Press.
Van Valin, Robert, and Randy LaPolla. 1997. *Syntax: Structure, Meaning, and Function*. Cambridge: Cambridge University Press.
Varela, Francisco, Evan Thompson, and Eleanor Rosch. 1991. *The Embodied Mind*. MIT Press.
Wertheimer, Max. 1923. Laws of organization in perceptual forms. Reprinted in W. D. Ellis (ed.), *A Source Book of Gestalt Psychology*, 71-88. London: Routledge and Kegan Paul.
Wittenberg, Eva, RJ, Gina Kuperberg, Martin Paczynski, Jesse Snedeker & Heike Wiese. in press. The processing and representation of light verb constructions. In A. Bachrach, I. Roy, and L. Stockall, (eds.). Structuring the Argument. John Benjamins.
Wittgenstein, Ludwig. 1953. *Philosophical Investigations*. Oxford: Blackwell.
Wynn, Karen. 1992. Addition and subtraction by human infants. *Nature* 358, 749-750.