To appear in W. Bechtel and G. Graham (eds.) *A Companion to Cognitive Science*, Basil Blackwell, 1997.

Structural Analysis

Robert Frank
March 4, 1997

The study of the human mind embodied in cognitive science strives to understand the nature of the abstract representations and computational processes responsible for our ability to reason, speak, perceive and interact with the world, etc. In addition, a commitment to a materialist resolution of the mind-body problem requires that we search for the manner in which these representations and processes are neurally instantiated in the brain. Given this dual aim, one might proceed in one of two ways: (1) bottom-upwards, commencing with the study of how low-level information and computations are encoded in the neuroanatomy of the brain, in the hope of working upwards toward an understanding of the properties of higher level cognitive processes; (2) from the top-downwards, using behavioral data from a variety of sources to provide an abstract characterization of cognitive processes, and then utilizing these results to guide our search for the neural mechanisms of cognition. At present, our understanding of the neural basis of higher level cognition is virtually negligible. Thus, it is probably fair to say that a contemporary study of any of the domains discussed in section II of this volume may only be carried out in a top-down fashion.

It is sometimes claimed that top-down research into cognition is incapable of producing anything but "castles in the sky" having little relevance to human thought. However, a brief look at the history of science shows us that the study of a phenomenon in the absence of an understanding of its underlying mechanisms is by no means novel and has moreover produced significant success. One prominent example is Gregor Mendel's study of heredity. By observing the external patterns in which properties of pea plants such as color and height were transmitted from one generation to the next, Mendel was able to deduce the existence of genes and the fundamental laws by which they combine. Of course, Mendel had no conception of the biological character of genes, nor any understanding of the reasons that his combinatorial laws held. Yet, his results have been substantially vindicated and formed the impetus for research into the biological basis of genetic material, culminating in Watson and Crick's discovery of DNA.

How might we mimic Mendel's methodology (and success) in the study of the mind? That is, how do we go about building an abstract theory of a cognitive process which can form the basis for subsequent neurological investigations? One route starts from the assumption that there are specialized representational structures which underlie mental processing within each cognitive domain. Just as Mendel's experience with pea plants led him to propose that heredity is best explained by positing an abstract representation, the gene, along with laws governing its behavior, so too can we use data from human behavior to lead us to the discovery of the abstract structures and laws governing cognition.

Within the cognitive sciences, this type of research has its roots within the domain of LANGUAGE, specifically within the area of *syntax*, the study of sentence structure. Any study of humans' capacity for natural language syntax must face the age old challenge of deriving infinite capacity (i.e., we can understand and produce arbitrarily long sentences, including those we have never before heard) from the finite resources provided by our physical endowment. During the first half of this century, work in the foundations of mathematics and computer science led to the development of a variety of mathematical and logical systems that fortuitously provided a formal means for answering this challenge. This led to innovative work by structuralist linguists such as Zelig

Harris which for the first time provided a set of mathematically precise representations and rules characterizing the well-formed utterances in a particular corpus of language use. Such precision represented a great advance, since it became possible for the first time to see the exact implications of analyses. Yet, structuralists largely adhered to a behaviorist stance on human psychology, and hence took their formal representations to be merely descriptive devices and of no psychological import. In the mid 1950's, Noam Chomsky broke with this behaviorist view, suggesting that the results of linguistic analysis should indeed be understood as the mental representations that comprise an individual's linguistic *competence*, i.e., the knowledge which underlies the ability to speak or understand a language. In this framework of generative grammar, Chomsky maintained the focus on providing a mathematically precise characterization of grammar, keeping some of the formal apparatus advocated by Harris while adding a new range of mathematical devices. However, he understood these devices as providing abstract descriptions of the mental computations underlying language (Chomsky 1986).

In the remainder of this article, we will have much to say on precisely how this use of mathematical structures has played a beneficial role in our understanding of human grammar. We will look at two case studies, each demonstrating that the computations involved in linguistic cognition attend to remarkably abstract representational structures, quite far removed from the stimuli of the external world. We will then briefly consider the mathematical foundations which support these linguistic investigations, and discuss the productive interactions that have resulted between studies of the properties of mathematical and grammatical structures. Finally, we will mention some other places in cognitive science where attempts have been made to make similar use of such abstract structures in modeling cognition.

Finding Structure in Language

Let us begin to investigate what kind of structures have been found to underlie human language. Consider the following English sentence:

(1) The student has finished her homework.

In order to make a question from this statement, we must change the order of the words, moving the auxiliary verb *has* to the front of the sentence:

(2) Has the student finished her homework?

We can ask the following question: what was the nature of the computation which affected this change in ordering? The simplest answer might go something like this:

(3) To make an English question, move the auxiliary verb to the front of the sentence.

This rule makes very few commitments about what structures underlie English sentences. It requires only that our linguistic computations recognize the notion *front of the sentence* and have the ability to identify elements in the category of auxiliary verbs, i.e., *have*, *be* or modal elements like *should*, *could*, etc. This simple formulation is, however, insufficient. In sentences involving two auxiliary verbs like that in (4), it does not tell us which auxiliary, *has* or *was*, to move.

(4) The student has finished her homework which was assigned today.

To address this problem, we need only complicate the rule slightly:

(5) Move the first auxiliary verb to the front of the sentence.

This rule makes only one additional ontological commitment, namely the notion of *first*. This notion has a straightforward translation into sensory terms, i.e., temporally earliest in the speech stream. Thus, it introduces no controversial assumptions about linguistic computation. Unfortunately, even this more complicated rule remains inadequate, as demonstrated by cases like the following:

(6) The student who is eating has finished her homework.

If we follow (5) and front the (temporally) first auxiliary verb, the element *is*, the result is a severely ill-formed string (where '*' indicates ungrammaticality) :

(7) * Is the student who eating has finished her homework?

To produce the well-formed version of this question, we must instead move the temporally second auxiliary, *has*:

(8) Has the student who is eating finished her homework?

It turns out that we must complicate our rule still further to achieve the desired result:

(9) Move the first auxiliary verb which follows the subject to the front of the sentence.

Since the string *the student who is eating* constitutes the subject, we move the next auxiliary verb in the sentence, i.e., *has*, producing (8). In the previous cases, the temporally first auxiliary is also the one which follows the subject. Observe that this rule differs from the previous one in that it makes reference to the abstract notion of *subject*, one which does not have any direct sensory characterization. Thus, we must assume that grammatical representations include a certain amount of structural analysis so as to allow the detection of the subject.

Recall that we take these grammatical rules and representations to form part of a speaker's knowledge of his or her language. Consequently, we are obliged to face the problem of how such rules and representations arise in the mind of a child during the process of language acquisition. From this perspective, it is interesting to note that in Crain and Nakayama's (1987) experiment, which ellicited sentences of the relevant type, none of the thirty 3- to 5-year-old English-learning children they tested ever gave a response like that in (7), as would be expected if they had fixed on the simpler, but incorrect rule for question formation in (5). This is quite puzzling in face of the observation that questions of the type in (8), which are necessary to distinguish this possibility from the correct one, are ordinarily (i.e., outside the experimental context) quite rare, so much so that it is not unlikely that they never occur in a child's linguistic input. Why, then, do children uniformly ignore the simpler possibility during their process of inducing the grammar of English, and instead proceed to hypothesize the more complex rule? Chomsky suggests that the resolution of this puzzle (often referred to as the *poverty of the stimulus*) lies in the recognition of a certain amount of innate grammatical knowledge, what he calls Universal Grammar (UG) (see WHAT ASPECTS OF THE COGNITIVE SYSTEM ARE INNATE?). Chomsky argues that UG predisposes children to learn grammars of a certain sort, in the case at hand, ones that utilize *structure dependent* rules, i.e., making reference to notions of abstract grammatical structure, such as subject, and not to other notions, even if definable directly in sensory terms. This structural dependence property of UG also explains why many apparently "simple" grammatical rules are absent from the languages of the world. For example, there is no language which forms its questions using rule (5), or by reversing the order of the words in the sentence, or by switching the second and fourth words. It's hard to find an independent reason why such things shouldn't be possible. Such rules are all simple to state and would be easy to compute, and a language which used them would be no worse off in terms of communicative possibilities. They simply do not seem to be part of any known human language. Reasons such as these make it quite difficult to provide functional explanations for the precise character of grammatical structures (see DOES COGNITIVE FUNCTION DETERMINE THE STRUCTURE OF COGNITIVE SYSTEMS?).

Let us turn to another type of evidence which brings the detail of grammatical representations into greater focus. Consider the following pair of sentences:

(10) a. John thinks he should eat spinach.
b. He thinks John should eat spinach.

The first of these has two possible interpretations: it is either a statement of John's beliefs about his own dietary demands, or alternatively John's beliefs about the dietary requirements of some other unidentified male. In contrast, the second may only be interpreted as a statement about the beliefs of someone other than John about John's own dietary needs. Similar contrasts hold in the following pairs:

(11) a. John's belief that he should eat spinach is unshakeable.
b. His belief that John should eat spinach is unshakeable.

(12) a. I gave John his spinach.
b. I gave him John's spinach.

In each pair, the pronoun may be understood as referring to, or *coreferential* with, John only in the first example and not in the second.

What is the explanation of these contrasts? As before, the simplest grammatical constraint ignores structural considerations and refers to the ordering among elements in the sentence:

(13) A pronoun may not be interpreted as identical in reference with a proper name if the pronoun precedes the proper name in the sentence.

Once again, however, the simple formulation is inadequate. In all of the following sentences, the pronoun precedes the proper name, yet they may both refer to the same individual, namely John.

(14) a. His mother thinks John should eat spinach.
b. After talking with his mother, John concluded I should eat spinach.
c. People who meet him in the morning find John irritable.

To understand the contrast between these cases and the (b) examples in in (10)–(12), we must consider some details of underlying grammatical structure. It is usually assumed that a sentence like (10b) has the following tree-like representation:

(15)

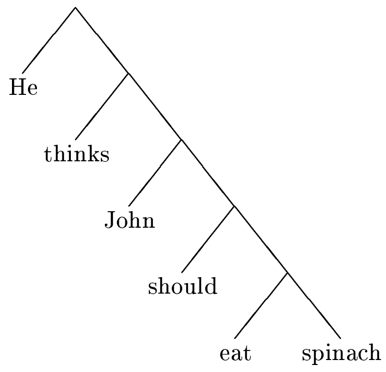

This structure represents the hierarchical groupings of the individual words into phrases, and these phrases into larger structures. For example, in this structure the verb *eat* forms a unit or *constituent* with the noun *spinach* since there is a single node in the tree which dominates them both but nothing else. In contrast, the pair *should eat* does not form a constituent since the only node which dominates them also includes the noun *spinach* (these three words together do, therefore, form a constituent). Constituency is explicable both in semantic and syntactic terms. We say the phrase, *eat spinach* is a constituent since it forms a semantic function, one which is predicated of the embedded subject *John*. From the syntactic side, we can now interpret our previous conclusion that grammatical operations are structure dependent as meaning that they may manipulate only constituents. From this, we may again conclude that the sequence *eat spinach* is a constituent since it may be moved to the front of the sentence (16) or deleted entirely (17):

(16) a. Eat spinach, he thinks John should.
b. * Should eat, he thinks John spinach.
(17) a. He thinks John should.
b. * He thinks John spinach.

What then is the structural relationship in (10b) which renders the coreference between *he* and *John* impossible? To help us see this, consider the structure of sentence (14a):

(18)

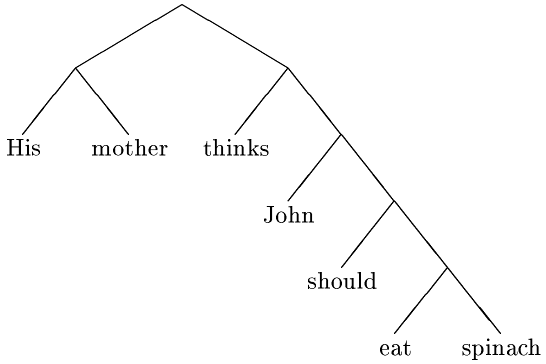

The sole difference between this structure and that above lies in the form of the subject constituent. In (15), this consists of the pronoun *he* alone. Here, however, the pronoun *his* is first combined with the noun *mother*, and this unit is then combined with the remainder of the structure. We can exploit this small structural difference to account for the interpretive differences we have already seen.

(19) A pronoun may not be interpreted as identical in reference with a proper name if the smallest (non-trivial) constituent containing the pronoun also includes the proper name.

By non-trivial constituent we mean one which includes something other than the pronoun itself. This constraint is quite natural from a structural perspective. It expresses in the simplest manner possible the structural relationship between the pronoun and the name *John* in (10b). What is striking is that our constraint in (19) accounts for a wide range of examples including all of those in (11) and (12) as well as those in (14). For example, given the following structure for example (14b), we can see why the pronoun may be co-referent with *John*:

(20)

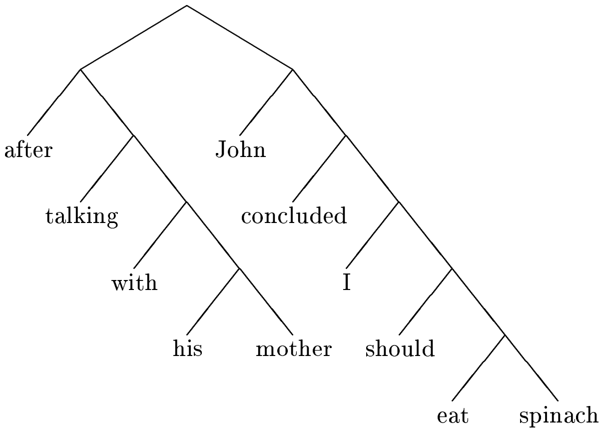

Since the smallest constituent containing the pronoun *his* is *his mother*, a constituent which does not include *John*, the constraint in (19) does not restrict the reference of the pronoun. Our success in providing an explanation for these facts about possible coreference provides compelling evidence both for the constraint in (19) as well as for the fine structure of our abstract grammatical representations to which it makes crucial reference. This constraint tells us something even more surprising about the structures underlying our grammars when we consider the following variation on example (14b):

(21) After talking with John's mother, he concluded I should eat spinach.

Once again, we find that coreference between *he* and *John* is blocked. Yet, our constraint does not seem to help us: The structure in this case is identical to that in (20), so that the smallest constituent containing the pronoun is *he concluded I should eat spinach*, which does not include *John.* Before despairing, let us take a look at the verb *talking* in this sentence. Although there is no visible (or audible) expression of the subject of this verb, our intuitions tell us that the subject, i.e., the "talker", is necessarily coreferent with the pronoun *he.* Traditional grammar accounts for this intuition by positing an invisible pronoun in this position, much like that which occurs in the subject of imperatives. If we take this traditional idea seriously and suppose that there is actually an invisible pronoun in the structure precisely where the subject would ordinarily appear, we can understand this case in terms of our constraint: since the smallest non-trivial constituent containing the invisible pronoun also includes *John*, *John* may not be coreferential with the invisible pronoun. As we observed above, this invisible pronoun is necessarily coreferential with the subject pronoun *he.* By transitivity, then, we derive the result that *John* may not be coreferential with *he.*

The morals of this second investigation are similar to those we drew above. We see again that it is necessary for grammatical computations to be sensitive to abstract properties of grammatical structure, including hierarchical relations and the presence of invisible pronouns. Furthermore, the constraint in (19) has interesting implications for the problem of the acquisition of linguistic knowledge. It has been shown experimentally by Crain and McKee (1985) that children as young as 3-years in age do not make errors of the sort that would result from ignorance of constraint (19), and it seems quite clear that they are never taught this constraint, either explicitly or implicitly. Furthermore, every known human language abides by constraint (19). This cross-linguistic invariance coupled with the poverty of the stimulus suggests that not only does UG innately specify the format for the linguistic rules which the child acquires, but also in some cases the rules themselves.

Choosing Structural Representations

In the last section, we saw how our understanding of human grammar has profited from structural analysis, i.e., the assumption that underlying our linguistic competence there are abstract rules and representations of a particular sort. Yet, our discussion of the properties of the grammatical rules and representations and why they were selected has been rather vague and informal. It is however clear that our success in applying structural analysis depends crucially upon choosing the appropriate underlying structures. Let us therefore turn briefly to the question of how this choice has been made in the case of grammatical analysis.

When applying structural analysis to some domain, we must make at least two choices. First, we must make the sometimes difficult decision of what objects are to serve as the atomic elements manipulated and combined in our abstract structures. Indeed, it is here that Mendel's contribution to the study of heredity lies, i.e., the recognition of the atomic notion of gene. In the case of syntax, the unit that we have assumed and that has proved appropriate for analysis is that of *word*. It is important to observe that the recognition of the word represents a significant abstraction from the observable data. No direct analog exists in spoken language: the speech stream contains no reliable indications of word boundaries, and two instances of what we might call the same word can vary significantly in how they sound. Nor is the notion of word reducible to anything related to units of meaning. Individual words can be composed of separate meaningful subparts (called *morphemes*) which are combined together in different words:

(22) a. help-less-ness
  b. help-ful
  c. care-less-ness
  d. care-ful

Instead, the notion of word comes to us through mankind's development of writing systems, many of which segment sentences into medium size chunks above the level of individual sounds. Note that our choice of words as the atomic unit of syntax is crucial: Our constraint on coreference, for example, would be obscured significantly if we were dealing with combinations of units smaller or larger than words.

Having identified our atomic units, we must address the second question of structural analysis: what kinds of rules and representations govern the combinations and manipulations of these atomic elements? In our two case studies, we have concluded that words are grouped into tree-like structures, parts of which may be transposed and/or deleted. Yet, our introduction of these abstract representations may have seemed a bit out of the blue. Why were such representations chosen? The answer to this question lies in a fortunate synergy in the first half of this century between studies of language and research into the foundations of mathematics and computer science. Around this time, formal models were being studied that could characterize the behavior of recently developed computing devices (see Partee et al. 1993 for an overview). These models provided a means of specifying an infinite number of objects/computational sequences using only a finite amount of information, much as a computer program can specify an infinite number of different outputs, depending upon the input. One such abstract model that was developed was that of *string rewriting systems*. In such a system, we start out with some basic string, usually the symbol 'S', and then iteratively rewrite portions of this string according to what the rules specify. Thus, given the system of rules in (23), we can carry out the sequence of rewriting steps in (24).

(23) $\mathrm{S} \longrightarrow \textit{frog}\ \mathrm{S}$
$\mathrm{S} \longrightarrow \textit{pond}$

(24) S ⟹ *frog* S ⟹ *frog frog* S ⟹ *frog frog frog* S ⟹ *frog frog frog pond*

Each rule tells us that the symbol to the left of the arrow ('S' in both rules in (23)) may be rewritten as the sequence of symbols to the right of the arrow. We continue with this process until there are no longer any symbols left which can be rewritten according to the rules. Notice that both of the rules in this system have an 'S' on their left, and therefore either may be chosen at any point in the process of rewriting. However, as soon as we choose the second rule, our rewriting terminates, since there will be no remaining occurrences of the symbol 'S'. It is for this reason that this sequence of rules allows us to generate all and only the strings which start off with a sequence of zero or more occurrences of *frog* followed by a single instance of *pond*.

As rewriting systems were developed, it became clear that they would also be useful for characterizing the sentences of human languages. Moreover, it was noticed that sequences of rewriting steps like that in (24) could be taken to generate a tree-like structure: at each point of rewriting, we add a layer in the tree, placing the newly inserted material below the rewritten symbol. Thus, the sequence of rewriting in (24) corresponds to the following tree:

(25)

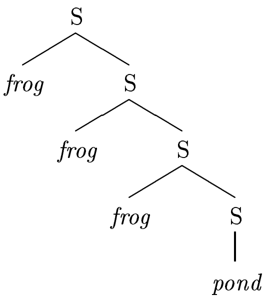

To produce the structures of natural languages, we can use systems of the same character. Thus, the following system of rules generates an infinite set of English sentences (assuming we replace the lower case part of speech symbols with an appropriate word):

(26) S ⟶ NounPhrase VerbPhrase
VerbPhrase ⟶ auxiliary VerbPhrase
VerbPhrase ⟶ verb NounPhrase
VerbPhrase ⟶ verb S
NounPhrase ⟶ name
NounPhrase ⟶ pronoun
NounPhrase ⟶ article noun

Interestingly, the sequences of rewriting steps which may be carried out using this system of rules correspond precisely to tree structures with the constituency we argued for above. The following is an example of one such structure:

(27)

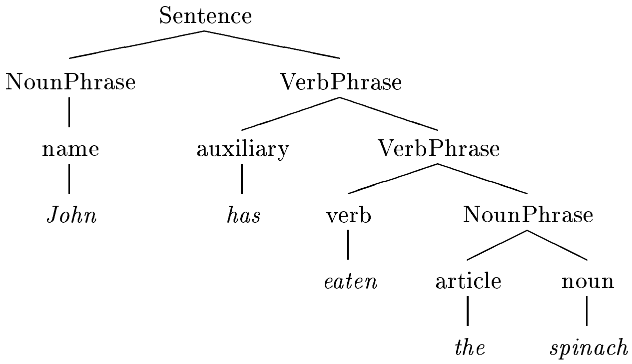

As the formal work produced a range of models which differed in their expressive power, linguists began to raise the question of whether the grammars of natural languages correspond in any interesting way to a particular class of mathematical devices. If this were true, it might allow us to understand restrictions on possible human grammars (embodied in UG) in formal terms. Chomsky (1957) argues that a certain class of devices, known as *finite state automata*, were too weak to serve as a formal basis for human grammars, since they cannot generate the recursive structural representations appropriate for certain sentences. Instead, he adopted the use of rewriting systems of the sort we just discussed, called *context-free grammars* (CFGs), to generate constituent structure representations. Chomsky also adopted certain operations, called *transformations*, which would manipulate the tree structures produced by the CFG, in order to handle cases like English question formation. It turned out however that the introduction of transformations expanded the expressive power of the system considerably. In response to this, a number of proposals have been made which aim to reduce the expressiveness of grammars by eliminating transformations from the grammar. It has however been shown that CFGs alone are insufficient for this because of the existence of certain constructions in Swiss German (Shieber 1985). This leaves open the question of the extent to which human grammatical systems may be formally characterized, though a number of promising possibilities exist (cf. Partee et al. (1993) for further discussion).

Abstract Structures in Other Domains

The methodology of structural analysis discussed in this article has been applied beyond the narrow realm of natural language syntax that we have discussed in this article. Within the study of language, similar methods of analysis have been pervasively applied to the study of sounds (phonology), words (morphology), and meanings (semantics), yielding a range of of abstract structural representations whose properties bear considerable explanatory burden. There are a wealth of cases in each of these domains analogous to those discussed here, though space prevents us from going in these (see Akmajian, Demers, Farmer and Harnish 1995 for a traditional overview, and Jackendoff 1994 for one more focused on connections with cognitive science). Additionally, these representations have shed substantial light on the processes of language acquisition and language change. It has been found that variation in the types of sentences that are used, whether during the course of children's acquisition of their native languages or in the centuries-long periods of linguistic change, are best characterized not as superficial and haphazard alterations, but rather in terms of parametric modifications to the fundamental underlying grammatical rules and constraints.

Moving outside the domain of language, one application of these same methods has been in the study of music cognition. Just as the representations of linguistic theory arise out of an attempt to model speakers' intuitions about well-formedness and possible meanings of the sentences of their language, Lehrdahl and Jackendoff (1983) present a cognitive theory of musical representations in an attempt to model the intuitions that listeners have concerning properties like meter and grouping of notes in a piece of music. Their representations have in common with the syntactic structures discussed above their hierarchical nature, but differ in at least two respects. First, as Lehrdahl and Jackendoff note, the constituents of syntactic structures are assigned a meaningful category, be it Sentence or Noun Phrase. Thus, what the first rule in (26) tells us is that the combination of a Noun Phrase and a Verb Phrase forms a Sentential constituent. In Lehrdahl and Jackendoff's musical structures, the combination of two constituents X and Y signifies only that one, say X, is the elaboration of the other, say Y. In this case, the resulting combination of X and Y is not essentially different in nature from the element Y occurring on its own. The second difference between the syntactic and musical representations lies in the mode in which their well-formedness is specified. We suggested above that syntactic representations are well-formed exactly when they accord with all of the rules and constraints of the grammar, like those in (9), (19) and (26). Lehrdahl and Jackendoff propose that well-formedness for musical structure is determined instead by a set of preference rules. These preference rules may at times conflict with one another, and in such cases, the conflict is resolved based on the strength of the relative preferences. Consequently, there will be some musical structures that violate some of the preference rules, yet are well-formed. Lehrdahl and Jackendoff point out that although their musical structures differ from syntactic representations, they are closely analogous to those exploited in another linguistic domain, namely that of phonology, particularly in the areas of stress and prosody, which are the rhythmic and melodic properties of spoken language. This is suggestive of a more general cognitive capacity of dealing with temporal organization of various sorts that could underly both musical and linguistic cognition. Concerning the different mode of specifying well-formedness, it is intriguing to note that the recently proposed *Optimality Theory* adopts the view that preference rules (or violable constraints) are indeed the appropriate means for characterizing the well-formedness of linguistic structure (Prince and Smolensky, to appear). If this view is on the right track, it suggests once more that the apparent differences across these domains might not be so large as they first seemed.

We also find work in visual PERCEPTION that exploits complex representations and constraints over them to explain behavioral data from early vision (Marr 1982), such as the "misperceptions" involved in visual illusions and subjective contours, as well as from high-level vision (Biederman 1987), like the ability to recognize objects in spite of variations in perspective or the presence of visual noise. To give one example, Biederman proposes that the basic three-dimensional appearance of objects is mentally represented as the combination of simple volumes called geons. The nature of the combinatorial relationships determines why some views of objects, such as rotations in the plane, are less easily recognized than others, like rotations in depth, as only the former induce a change in the TOP-OF relation and thereby present a visual stimulus which does not exactly match the mental representation.

Finally, results from a number of studies of REASONING and PROBLEM SOLVING have been used to argued that humans use structured representations in carrying out these tasks (Johnson-Laird 1983, Newell and Simon 1972). Johnson-Laird, for instance, attempts to explain the difficulties people experience in performing certain types of deductive inference. He suggests that people do not perform inferences from statements like "All scientists are sceptics" and "Anne is a scientist" to "Anne is a skeptic" by manipulating the standard sorts of logical inference rules, but instead by constructing mental models of the situations which are described. Some collections of statements do not uniquely describe one single mental model and an individual who is to carry out correct inference must simultaneously consider multiple models. Johnson-Laird demonstrates that the level of difficulty for inferences is predicted by the number of mental models that must be considered (as well as the order in which they must be examined), rather than by the character or length of steps of application of logical inference rules.

In these latter domains, abstract mathematical structures have been important in certain lines of research, particularly research focused on computational issues. Yet, in work tightly linked to empirical investigations, structures of the type brought up in our discussion of syntax have not played the significant role in shaping research directions and theories that it has in the study of language. One might interpret this state of affairs as demonstrating that the abstract structures appropriate for these other domains differ quite radically from those used in linguistic cognition. Indeed, it may well be that such structures are of an as yet undiscovered character. If this is true, this might be an indication that different cognitive domains are processed in distinct modules, and hence do not share the same types of representations (see IS THE COGNITIVE SYSTEM MODULAR?). An alternative explanation of the difference in the success of structural analysis across different domains might tie it to our not yet having found the appropriate atomic elements which the abstract representations combine and manipulate. In the case of high-level vision, for example, determining the right set of primitives from which object representations are composed remains a matter of significant debate. Note that the problem of atomic units is significantly more difficult in visual perception and reasoning than language, largely due to the relative paucity in external evidence for the nature of internal representations. If this is the correct diagnosis, we can hope that it is simply a matter of time before more of the cognitive sciences are able to unlock the underlying laws and principles of cognition that the mind still holds secret.

Bibliography

Akmajian, A., Demers, R.A., Farmer, A.K., and Harnish, R.M.: *Linguistics, An Introduction to Language and Communication* (Cambridge, MA: MIT Press, 1995).

Biederman, I.: ‘Visual object recognition’, *Visual Cognition: An Introduction to Cognitive Science*, ed. S.M. Kosslyn and D.N. Osherson (Cambridge, MA: MIT Press, 1995), pp. 121–165.

Chomsky, N.: *Syntactic Structure* (The Hague: Mouton, 1957).

Chomsky, N.: *Knowledge of Language, Its Nature, Origin, and Use* (New York: Praeger, 1986).

Crain, S. and McKee, C.: ‘Acquisition of structural restrictions on anaphora’, *Proceedings of the 16th North Eastern Linguistics Society Meeting*, ed. S. Berman, J.-W. Choe, and J. McDonough (Amherst, MA: Graduate Linguistics Student Association), pp. 94–110.

Crain, S. and Nakayama, M.: ‘Structure dependence in grammar formation’, *Language* 63 (1987), 522–543.

Jackendoff, R.: *Patterns in the Mind: Language and Human Nature* (New York: Basic Books, 1994).

Johnson-Laird, P.: *Mental Models* (Cambridge, MA: Harvard University Press, 1983).

Lehrdahl, F. and Jackendoff, R.: *A Generative Grammar of Tonal Music* (Cambridge, MA: MIT Press, 1983).

Marr, D.: *Vision* (San Francisco: Freeman, 1982).

Newell, A. and Simon, H.A.: *Human Problem Solving* (Englewood Cliffs, NJ: Prentice Hall, 1972).

Partee, B., ter Meulen, A., and Wall, R.E.: *Mathematical Methods in Linguistics* (Dordrecht: Kluwer, 1993).

Prince, A. and Smolensky, P.: *Optimality Theory: Constraint Interaction in Generative Grammar* (Cambridge, MA: MIT Press, to appear).

Shieber, S.: 'Evidence against the context-freeness of natural language', *Linguistics and Philosophy* 8 (1985), 333–343.