A Suggestion For A Linguistics With Connectionist Foundations

George Lakoff
Department of Linguistics
University of California
Berkeley, CA 94720

Abstract

The theory of cognitive linguistics (as outlined in my book *Women, Fire, and Dangerous Things* and Langacker's *Foundations of Cognitive Grammar*) converges with connectionist cognitive science in a variety of ways. This paper gives an overview of what those convergences are and what an overall theory of language with connectionist foundations might look like.

1. Some Larger Issues

I would like to situate this lecture with respect to certain larger issues raised by connectionist research. I see connectionist research as being primarily concerned with the question:

> How is it possible for the brain to *be* the mind? That is, how is it possible for a physical brain to engage in reason? And how is language represented in the brain?

Such questions are indirectly related to the question:

> How could human intelligence have evolved from the brains of other primates?
>
> Does human intelligence involve the use of capacities present in the animals from which we evolved?

I will be presenting a number of linguistic results that bear indirectly on these issues, and will add a number of speculations that, if correct, might allow us to begin answering such questions. The linguistic results I will be presenting indicate that human reason uses some of the same mechanisms involved in perception and that human reason can be seen as growing out of perceptual and motor mechanisms.

Before going on to this, however, it is important to make a distinction between connectionist modelling and the connectionist theory of mind.

2. Activation Patterns Over Portions of Topographic Maps Are Meaningful

The various parts of the body are connected to the brain by neural networks. Indeed, they are literally 'mapped' onto sheets of neurons in the brain, called 'topographic maps'. Such mappings preserve topological relations, but not relative size; e.g., the topographic map of the thumb is next to that of the forefinger, but is much larger than it.

The neural networks of the brain are an appropriate locus for a theory of how concepts are embodied for a very simple reason:

> Patterns of activation over topographic maps in the brain are meaningful, because of the way those portions of the brain are connected to the body.

Consider, for example, the pattern of activation that arises in those neurons in the topographic map of the retina that characterize color. The patterns of activation over such neurons characterize color categories, and we experience those patterns as colors within those categories. Those same patterns of activation, located in some other topographic map in the brain, would be experienced as something else: say, in the map of auditory space, they would be experienced as sounds. Activation patterns considered in isolation from their locations in the brain are meaningless in themselves; at particular locations in the brain, they are meaningful — not by virtue of the location per se but because of the way such locations are connected to the body.

It is in this sense that the neural networks of the brain (as it is hooked up to the body) appear to be the right locus for a theory of how concepts are embodied. They are thus a suitable basis for a theory of meaning, given the additional assumption that our perceptual systems respond in a lawlike way to external stimuli.

3. The Connectionist Theory of Mind versus Connectionist Modelling

Connectionist cognitive science, as I understand it, is ultimately about the physical brain as located in the body. It is important to see how connectionist modelling fits into the overall program of connectionist cognitive science. Models of neural networks, simulated on a computer, are very different from the physical brain as it is located in the human body. Computer models of neural networks are disembodied. Their job is limited. Their job is to study the properties of particular neural architectures,

of data. Such models have been providing us with important insights, and they are vital to the discovery of the general properties of neural networks. But such computer models, being disembodied, can never by themselves show how meaning is embodied. An account of the relevant neurobiology will be required as well, in addition to an account of how linguistic generalizations are represented in neural networks.

As a cognitive scientist and a linguist, I am interested primarily in connectionism as a crucial part of the study of how the physical brain can be the mind — in particular, how the properties of human language and human reason can be seen as emerging from the properties of our embodied neural architecture. I see my role as helping to provide the linguistics necessary to make this a reality. I will begin by discussing two kinds of results in cognitive semantics that seem to bring that reality closer.

4. Cognitive Topology

Consider the following problem: Imagine watching a game of ping pong. Each time the ball is hit, it moves with a different trajectory with respect to the net. There are, of course, a potential infinity of such trajectories. Yet, as observers and speakers of English, we can classify them into a very small number of categories. The ball either goes *over*, *under*, *into*, *through*, or *around* the net. Each preposition thus defines an infinite category of scenes. Somehow, we have the ability to correctly categorize an infinity of such scenes. How can we do it? How can we represent the meanings of those prepositions so that we can answer this question? What kind of concepts permit such an infinite categorization of visual scenes?

The answer given by cognitive linguists is that the spatial meanings of English prepositions are 'topological' in character, topological in the sense that they generalize over geometry. Consider, for example, the central sense of *over*. It breaks down in the following way: There are two BOUNDED REGIONS (one on one side of the net, one on the other). There is a PATH (from one BOUNDED REGION to the other). There is an OBSTACLE (the net) and a VERTICAL ORIENTATION to that obstacle. The PATH goes through a region on the vertical axis above the obstacle. There is a lack of CONTACT between the PATH and the OBSTACLE.

Notions like BOUNDED REGION (sometimes called CONTAINER), PATH, OBSTACLE, CONTACT, and VERTICAL ORIENTATION are a set of recurrent structures that appear throughout the languages of the world. These structures are topological in the sense that they apply to, but generalize over, particular geometries. They are also oriented relative to certain orientational parameters defined by the human body. There are a relatively small number of them, and they recur in the senses of many morphemes of English and other languages. For example, the central meanings of *in* and *out* use the BOUNDED REGION schema, the central meaning of *across* uses the PATH schema, and so on.

One of the most interesting properties of image-schemas is that they have built-in logics. For example, BOUNDED REGION schemas, also called CONTAINER-schemas, have essentially a Boolean logic. Consider two CONTAINER-schemas, A and B, such that A is in B, and an object X is in A. (See Figure 1.) We 'see' instantly, without doing any logical deduction, that X is in B. All we need to do is shift our focus to the relationship between X and B.

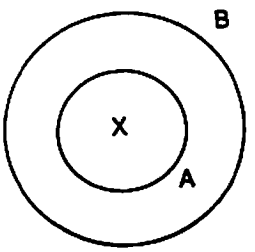

Figure 1: Containers are bounded regions.

Let us take another example. Consider a PATH P, and an object X that travelled up to a certain point on that path. We know that it has been at all previous points and has not been at any subsequent points.

In general, image-schemas have logics built into their topological structures, and spatial inferences arise via the application of attentional mechanisms.

5. Metaphor

Metaphor used to be thought of as a poetic use of language. Over the past decade it has become clear that poetic metaphorical expressions arise by virtue of a general type of cognitive mapping mechanism that is used throughout ordinary natural language semantics as a means to understand abstract concepts in terms of more concrete ones.

An early example of such a mapping was discovered when a Berkeley student was told by her boyfriend that their relationship had hit "a dead-end street." Analysis revealed that this was part of a systematic way of comprehending love in terms of a journey, with many conventional expressions of this metaphor in ordinary English:

Look *how far we've come.*
We're *at a crossroads.*
Their marriage is *on the rocks.*
We're *spinning our wheels.*
This relationship isn't *getting anywhere.*
It's *bogged down.*

We're *drifting apart.*

We'll have to *go our separate ways.*

It's been a *long, bumpy road.*

The metaphor does not reside in any of these particular expressions.

Rather there is a metaphor way of understanding love in terms of a journey by the following conceptual mapping from the spatial domain of journeys to the emotional domain of love:

TRAVELLERS correspond to LOVERS

THE VEHICLE corresponds to THE LOVE RELATIONSHIP

PHYSICAL CLOSENESS corresponds to INTIMACY

DESTINATIONS correspond to SHARED GOALS

THE PATH corresponds to THE COURSE OF THE RELATIONSHIP

IMPEDIMENTS TO TRAVEL correspond to DIFFICULTIES

This metaphorical mapping has entailments. For example, if the relationship isn't getting anywhere, then you have either be satisfied with not achieving shared goals, or you have to fix up the relationship, or you have to abandon it.

Examples of conceptual mappings of this sort have been discovered by the hundreds, and it has become clear that many of the most fundamental concepts in our conceptual system are understood via such metaphorical mappings from concrete to abstract domains. Consider, for example, the EVENT STRUCTURE METAPHOR, in which events are understood in terms of motion in space:

LOCATIONS correspond to STATES

MOVEMENTS TO NEW LOCATIONS correspond to CHANGES TO NEW STATES

SELF-PROPELLED MOTIONS correspond to ACTIONS

FORCES RESULTING IN MOTION correspond to CAUSES

DESTINATIONS correspond to GOALS

PATHS TO DESTINATIONS correspond to MEANS

This metaphorical mapping has a great many entailments which are realized in hundreds of everyday expressions. For example, IMPEDIMENTS TO ACTION are understood in terms of IMPEDIMENTS TO MOTION:

We're at an impasse.

I've hit a brick wall.

We're in rough waters.

Similarly, EASE OF ACTION is understood in terms of EASE OF MOTION:

It's all downhill from here.

It's smooth sailing from here on.

FORCED ACTION is understood in terms of FORCED MOTION:

He pushes me too hard.

I was dragged into doing it.

My boss really throws his weight around.

She leads him around by the nose.

PROGRESS in acheivng some goal is MOVEMENT TOWARD A DESTINATION:

We're moving ahead on this project.

There's nothing standing in our way.

We only have a short way to go.

We can see the light at the end of the tunnel.

I'm weighed down with a lot of other projects.

We're sliding backward.

He's floating aimlessly.

He needs some direction.

He's getting nowhere with this project.

One of the most interesting findings in metaphor research is that metaphors can map cognitive topological structures from spatial to nonspatial domains. Classical categories are understood metaphorically in terms of CONTAINER-schemas. It is this understanding of categories as containers for their members that makes the logic of bounded regions into Boolean logic. Thus, we saw that if CONTAINER A is in CONTAINER B, and X is in A, then X is in B. Under the CATEGORIES-as-CONTAINERS metaphor, this spatial reasoning has the effect of the classical syllogism: If All men are mortal and Socrates is a man, then Socrates is mortal. In mapping CONTAINER-schemas from the spatial domain to the domain of categories, the CATEGORIES-as-CONTAINERS metaphor preserves the topological structure of the CONTAINER-schemas and with it, the internal logic. Using cognitive topology plus metaphor, the effect of classical syllogisms is achieved with with no deduction, that is, with no maniuplation of symbols. All that is involved are (1) the cognitive topological structure of the CONTAINER-schema, (2) the metaphorical mapping, and (3) the ability to shift focus.

Let us take another example. We saw above that there is a PATH-logic inherent in the structure of paths. Thus, if you are moving along a path and are located at point X, then you have been located at all points before X and not at any points following X. There is a metaphor that maps paths into linear quantity scales, as the following expressions indicate:

John had *far* more money than Bill.

John's wealth *goes beyond* the imagination.

John is *way ahead* of Bill in intelligence.

In a variety of languages, the word for "more than" is the word whose central meaning is "pass" (cf. the origin of "surpass" in English).

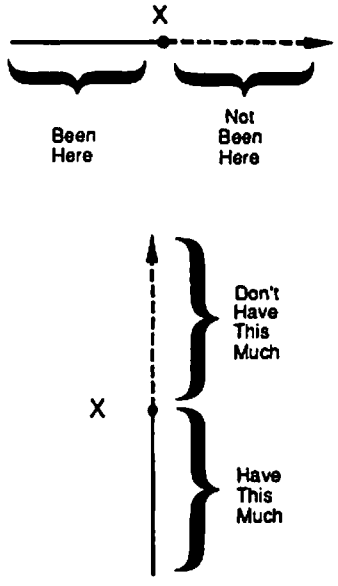

Figure 2: Linear scales resemble paths.

By virtue of this metaphor, the spatial logic of paths becomes the logic of linear scales, as in Figure 2. Previous points on the path are mapped into lower amounts, and subsequent points into higher amounts. Take the inference that, if you have reached point X on a path, then you have been at previous points. This gets mapped by the metaphor into the inference that if you have X amount of money in the bank then you have all smaller amounts there. For example, if you have $500, then you have $400, $300, $200, and so on. Similarly, the inference that if you have reached point X, you have not reached subsequent points gets mapped into the inference that if you have exactly X amount of money, then you do not have more. Again, all this can be done without any deduction. It is a consequence of the internal structure of the PATH-schema, plus the metaphor, plus the ability to shift focus.

Cognitive topology is a mechanism by which we impose structure on space, in such a way as to give rise to spatial inferences. Metaphor is a mechanism for mapping that structure onto nonspatial domains, preserving the cognitive topology and hence the inferential structure. Jointly, cognitive topology plus metaphor provide a way of linking the visual system to abstract reasoning. The linguistic results in these areas suggest that human reason is not deductive in character, that is, not a matter of symbol manipulation, as in predicate calculus. Rather, it is topological in character, yielding inferences by means of attention shift within a complex topological structure.

What all this suggests is that our conceptual system makes use of mechanisms that arise in the sensorimotor system. Fodor and Pylyshyn have argued that the only known way to account for human reason is by symbol manipulation in the manner of deductive logic. Cognitive linguistics suggests that human reason is not like that at all, and that it uses mechanisms that do not involve predicate-calculus-like symbol manipulation.

6. Implementation: Regier's Conjecture

Cognitive Topology and Metaphor fit well with the Connectionist Theory of Mind. The first connectionist implementation of a nontrivial portion of cognitive topology was done by Terry Regier, a graduate student at Berkeley, as his project for the 1988 Connectionist Summer School (see his paper in this volume). Regier made use of a connectionist implementation of a version of Shimon Ullman's visual routines in characterizing image-schemas. He was even able to characterize two of Brugman's image-schema transformations using threshold functions. On the basis of work done so far, Regier has conjectured (personal communication) that Ullman-style visual routines (which he has reworked in connectionist terms) are sufficient to characterize all known structures in cognitive topology.

I do not know if Regier's conjecture will stand up to rigorous empirical testing. First, it must be applied to the inventory of known image-schemas. Second, further research on image-schemas, both in English and other languages must be done, so that Regier's techniques can be applied to a full range of elementary structures. Preliminary research indicates that there are not all that many elementary cognitive topological structures used in the world's languages. Still, this must be established as well as possible on the basis of data from a wide variety of language families.

Regier's conjecture, of course, is dependent on some version of Ullman's visual routines. Since the validity of Ullman's general approach is an empirical matter, Regier's approach to the implementation of cognitive topology will falter if Ullman's approach does.

Nonetheless, Regier's work represents a significant start on the connectionist implementation of cognitive topology, and hence on the linking of the visual system with the conceptual system. What Regier has done is to show that certain nontrivial aspects of cognitive topology are implementable. Moreover, Regier's Conjecture defines a significant long-term reserach program on the relationship between abstract reasoning and the visual system.

Let us now turn to metaphor. On the basis of research done by Mark Turner and myself (Lakoff and Turner, forthcoming), it is our conjecture that

metaphorical mappings in general project the cognitive topology of the source domain on the target domain. There is no room here to go into our reasons for that conjecture. If we are right, then a general mechanism for the implementation of metaphorical mappings may be at hand, even for cases of mapping across a narrow channel. Gary Cottrell has constructed a back-propagation network that can learn to map pictures across a narrow channel. Networks with such a capability ought to be capable of projecting the cognitive topology of one domain onto another. This has not yet been tested, but given Regier's implementation of cognitive topology, it should be testable before too long.

7. Some Convergences Between Cognitive Linguistics And PDP Connectionism

There are a number of ways in which PDP connectionism and cognitive linguistics converge. The convergence is defined relative to a basic metaphor that links the two fields.

The Basic Metaphor Linking Cognitive Linguistics to Connectionism: a linguistic pattern (in CL) is a pattern of connection strengths (in PDP).

This metaphor is, at present, vague. To be made precise it will require a way to map precisely from PDP to CL. If the mapping were fully precise, it would enable us to understand linguistic principles in connectionist terms. The metaphor is, however, sufficiently precise for us to compare certain of the basic properties of PDP and CL.

Generalizations

PDP: Patterns of connection strengths constitute the generalizations that emerge in neural networks.

CL: Linguistic patterns constitute the generalizations characterizing knowledge of a language.

In a PDP implementation of cognitive linguistics, each PDP generalization would correspond to a CL generalization.

Constraint Satisfaction Systems

PDP: PDP systems are constraint satifaction systems.

CL: Cognitive grammars are constraint satisfaction systems.

Meaningfulness

PDP: Activation patterns at fixed locations in the brain are meaningful (they are not like arbitrary symbols that have to be given a meaning).

CL: Structures in cognitive topology and basic-level concepts (Lakoff, 1987, Ch.2) are meaningful in themselves (they are not like arbitrary symbols that have to be given a meaning).

In both cases, the patterns characterizing linguistic regularities are meaningful in themselves (whereas in a symbol-system, the symbols are meaningless in themselves and have to be given meanings).

8. Sensorimotor Grounding For Semantics And Phonetics:

PDP: Semantics and phonetics are both grounded in the sensorimotor system; they are thus autonomous, in the sense that they are independent of syntax and morphology. Syntax and morphology involve patterns of correlations between semantics and phonology.

CL: Semantics and phonetics are cognitively "autonomous"; syntax and morphology involve conventionalized patterns of correlations between semantics and phonology.

The idea here is to explain how language might have evolved and how it might be learned without hypothesizing a magical "language organ" with no sensorimotor basis. CL and PDP contrast sharply here with generative grammar. Generative grammar is characterized mathematically in terms of recursive function theory, that is, in terms of a symbol manipulation system that cannot make use of any interpretation of the symbols. Thus, an autonomous syntax with generative grammar must be a module that cannot be affected by any general cognitive or sensorimotor input. Real modules in the brain are, of course, not like that. They all have inputs and are affected by them, and they are all ultimately linked to the sensorimotor system.

In the PDP Connectionist Theory of Mind at present, the idea of an autonomous generative syntax does not make any sense, since syntactic categories and structures would not have any grounding of any sort. For example, what would it mean for a pattern of connection strengths in some part of the brain to characterize ungrounded purely abstract notions like X-bar as opposed to some other arbitrary concept? The very idea would be senseless in a PDP Theory of Mind.

Symbolization

PDP: Symbolization occurs when there are appropriate connections linking a semantically grounded pattern of connection strengths with a phonetically grounded pattern of connection strengths.

CL: Symbolization occurs when there is a pairing of semantic and phonological elements.

In natural languages, symbolization has to do with the relation been some concept to be expressed and some means of expression (in sound, gesture, writing, etc.). Such a notion of symbolization is natural in PDP connectionism.

Nonfinitary nature

PDP: Activation patterns are continuous, not finitary.

CL: Structuring is topological, not finitary.

The topological structuring used in cognitive linguistics allows for both boundedness and inferential structure within nonfinitary structures. Such topological structuring can be implemented in PDP systems. This is extremely

important because language does show a large amount of inferential structure. Cognitive topology can characterize that structure in a nonfinitary way, and thus permits some of the most common criticisms by generative linguists to be met.

Constructions

PDP: Activation patterns can exist across neural ensembles connected by narrow channels.

CL: "Constructions" = multi-dimensional schemas with connections across the dimensions.

A construction is a linguistic pattern characterizing a linguistic generalization. A construction can be purely semantic or purely phonological. Syntactic and morphological constructions (in CL) are multi-dimensional; they link aspects of semantics with aspects of phonology.

A given neural ensemble can be seen as characterizing a dimension of linguistic structure:

- A pattern of connection strengths within an ensemble corresponds to a schema within that dimension of structure.
- Connections across ensembles correspond to links across dimensions of structure.
- The entire pattern (including both within-ensemble and cross-ensemble patterns of connection strengths) corresponds to the construction as a whole.

A grammar in CL is a (highly structured) collection of constructions. A grammar in PDP is a collection of patterns of connection strengths of the appropriate sort.

Compositionality

PDP: Composition of activation patterns by superposition.

CL: Composition of constructions by superposition.

Incidentally, the kind of compositionality that PDP and CL share, namely, the superposition of patterns of a nonfinitary nature is not the kind of compositionality that occurs in formal grammars.

Sound symbolism

PDP: Predicts the naturalness of sound-symbolism (as patterns of correlations between semantics and phonology).

CL: Sound-symbolism exists in the form of patterns correlating semantics and phonology.

9. Structure In The Grammar And In The Lexicon:

PDP: Schema variability.

CL: Radial Categories.

Constructions (whether lexical or grammatical in nature) have been found to form radial categories, that is, networks consisting of a central construction and related noncentral constructions. The noncentral constructions are variants of the more central constructions in such a network. This characteristic of constructions corresponds to a characteristic of schemas within PDP systems, namely, schema variability (cf. Rumelhart & McClelland, Ch. 14). Schema variability in PDP also appears to correspond to the concept of "motivation" in CL: If schema B is a variant of A, then A can be seen as motivating B.

Inheritance of properties

PDP: If schema A is a variant of schema B, A will automatically inherit all the properties of B except for those that make it distinct from B.

CL: Peripheral schemas in a radial network inherit properties from more central schemas.

Degree phenomena

PDP: Patterns may match inputs to a degree.

CL: Degrees of well-formedness exist in grammar.

Gradual Grammaticization

PDP: Patterns of connection strengths arise gradually.

CL: Constructions enter the language gradually, and at any one time, constructions or lexical expressions may be conventionalized only to a degree.

Incidentally, gradual grammaticization does not make sense in generative grammars, where the rules have to be either present or not.

General And Special Cases

PDP: Generalizations involve low activation levels over large regions of the network; special cases involve higher activations over small regions of the network, and these override the low activations of the general cases.

CL: Panini's Law: Special cases preempt the general case.

Panini's Law falls out as a general property of PDP systems.

Grounding of Phonetics

PDP: Regions in the acoustic-articulatory phase space.

CL: Phones.

Phones are to be characterized in terms of correlations between acoustics and articulation. The acoustic characteristics of a phone and its corresponding articulation would, from a connectionist point of view, both be characterized in terms of vectors in an articulatory-acoustic phase space. The vectors represent activation patterns over the acoustic and articulatory topographic maps.

Grounding of Basic-level Concepts

PDP: Regions in a multi-dimensional sensori-motor phase space.

CL: Basic-level concepts.

Basic-level concepts are characterized by correlations among gestalt perceptions, motor programs, and imaging capacity. Each of these, presumably, would be representable in terms of regions in a sensori-motor phase space, and their correlations would then be representable in a multi-dimensional phase space.

Grounding of Cognitive Topology

PDP: Units governing certain type of sensori-motor coordination (e.g., interacting with containers, following moving objects, and so on).

CL: Cognitive topology: bounded regions, paths, links, various orientations, etc.

If Regier is correct, cognitive topological structures can be characterized using a connectionist version of Ullman-style visual routines.

Metaphor

CL: The projection of the cognitive topology of one conceptual domain onto another conceptual domain.

PDP: The kinds of analogical structure-mappings of the kind used in metaphor are natural in PDP systems.

In CL, metaphor is the main mechanism for accounting for abstract concepts: they are understood in terms of more concrete concepts. The mechanism is the projection of the cognitive topology of a source domain onto a target domain. Since the cognitive topology characterizes the inferential structure of the source domain, the inferential structure is mapped across.

The source and target domains are seen as neural ensembles connected by a small number of connections (a narrow channel). The structure of the source domain concept is characterized by a schema. With the right architecture, a source domain pattern can be projected through the narrow channel to match or create a target domain pattern.

Metonymy

PDP: Activation of one part of a net resulting in the activation of another part (or of the whole net).

CL: Metonym: A mapping from one element of a schema to another element (or to the whole schema).

Networks that do metonymy are now being developed by Dave Touretzky at Carnegie-Mellon.

Where This Leaves Us

The Basic Metaphor of Connectionist Linguistics — Linguistic Patterns Are Patterns of Connection Strengths — allows us to see convergences between CL and PDP. If the metaphor were elaborated into a fully explicit mapping, along the lines suggested, it would allow us to comprehend linguistic phenomena in connectionist terms. If the mapping were fully explicit, the "convergences" could be turned into explanations, in which characteristics of language would be explained in terms of properties of neural networks. Such an explicit mapping would also provide characterizations of such notions as "basic-level concept", "cognitive topology", "metaphor", "metonymy", "construction", "grammar", and so on.

Given such basic notions, it would be the job of cognitive linguistics to show how language in its full complexity could be built up. For example, we would have show how such notions could combine to form Idealized Cognitive Models, Mental Spaces, Categories (with their full range of types of prototype phenomena), as well as such notions as Causation, Aspect (that is, event structure), Purposes, and so on. We would also have to show how to characterize semantic roles (for instance, Agent, Patient, Goal), grammatical categories (Noun, Verb, Adjective, Preposition) grammatical relations (Subject, Direct Object), long-distance dependencies, complex grammatical constructions, and the full range of syntactic phenomena. Let us now consider how this might be done.

10. Cognitive Linguistics With Connectionist Foundations

Here are the basic ingredients for a cognitive linguistic theory:

- Cognitive Topology
- Metaphor
- Metonymy
- Basic-level concepts
- Phones
- Constructions
- Composition by superposition

It is the job of cognitive linguistics to show how these ingredients combine to form conceptual systems and languages based on them. At the outset, one would have to be able to show how the following basic constructs of cognitive linguistics could be formed.

- ICMs
- Categories (with sources of prototype effects)
- Abstract concepts and abstract inference patterns
- Mental spaces
- Semantic Roles
- Grammatical Categories
- Grammatical Relations
- Syntactic and phonological constructions

Here are some general ideas for how to go about it.

ICMs

Start with cognitive topology; each image-schema has an internal logic with spatial inference patterns. Add metaphors to map them into abstract domains, to yield temporal order, event structure, causation, purposes, category structure, mental space structure, etc.

Yield: General inference patterns.

Given the basic superstructure of ICMs, you add basic-level concepts to ground the category system.

Yield: ICMs.

Categories

Each category is a metaphorical container.

Fuzzy categories: apply dimension shift to boundary, making it one-dimensional and superimpose on it a linear scale.

Classical categories within ICMs: Identify elements of ICMs as having metaphorical container structure.

Metonymic Categories: Allow an ICM characterizing a member or subcategory to stand metonymically for another member, subcategory, or the whole category, for the purpose of drawing some particular kind of inference.

Radial Categories: Take a metaphorical container. Superimpose a center-periphery schema. Place the central subcategory at the center. Add links linking more central to less central categories, where the links are characterized by metaphor, metonymy, ICM-addition, image-schema transformation, etc.

Grammatical Categories

Grammatical categories are formed using semantic categories and mechanisms of radial category construction.

Noun: The center consists of single bounded regions in physical space: objects, beings, locations. Extensions include metaphorical objects, beings, and locations: e.g., times, ideas, etc.

Verb: The center consists of all processes, that is, motions and metaphorical motions.

Adjective: The center consists of properties of objects (e.g., color, size, etc.)

Preposition: These are topological concepts, with no basic-level content. The central cases are in the spatial domain.

Grammatical relations

Subject: A radial category, whose center is both agent and topic.

Direct Object: A radial category whose center is patient.

In CL, grammatical categories and relations are, thus, neither primitive, nor autonomous. They are built up using semantic categories and mechanisms of category-formation (e.g., radial categorization). This allows them to be characterized ultimately in terms of those things that are autonomous.

Grammatical Constructions

These pair ICMs characterizing the semantics and pragmatics of constructions with information about grammatical categories, grammatical relations, and various aspects of phonology: the phonological forms of grammatical morphemes, aspects of intonation, and constraints on the order in which the elements are to be spoken. (For examples, see Lakoff, 1987, case study 3.)

Phonological Constructions

Phonological constructions are patterns of correspondences across three dimensions of structure:

(M) Morpheme structure: Individual phonological forms stored in memory.

(W) Word structure: Assemblages of morphemes in structured sequences of words.

(P) Phonetic Structure: Forms as they are to be uttered.

Each phonological construction is either (1), a constraint within one of these dimensions, e.g., a constraint on morpheme structure; or (2) a constraint across two of these dimensions. Phonological constructions combine by superposition; there is no such thing as rule ordering or derivations.

Before we go on to consider an example in detail, it is important to compare constructional and generative phonology with respect to implementation in neural networks. Generative phonology contains derivational sequences which are ordered atemporally — outside the flow of real time. To my knowledge, such atemporal derivations cannot be implemented in real (biological) neural networks, which operate only in real time. If this is correct, as it seems to be, generative phonology appears not to be implementable in the brain.

On the other hand, multiple activation patterns can be activated simultaneously in real time in real neural networks. Phonological constructions should be representable in terms of such patterns. Constructional phonology is, thus, an attempt to characterize phonological regularities in a language in terms of structures capable of being implemented in neural networks.

Unfortunately, much of modern phonology has been formulated in generative terms and thus will have to be rethought. The following example is presented in the spirit of beginning such a rethinking of phonology.

The first, and most obvious, problem facing such a rethinking is that generative phonologists have given arguments and evidence in favor of rule ordering. What needs to be shown, in the face of these arguments, is that they rest on the formal generative framework assumed,

and disappear when a constructional framework is introduced. The example to follow is presented for that purpose. It is not intended as an argument that constructional phonology is superior on purely linguistic grounds. I think such arguments can be given, but they would go beyond the scope of this paper.

The example given is from Mohawk. It is a textbook example taken from Halle and Clements' *Problem Book in Phonology*, pp. 121-123. In that book, this problem is intended to show that complex rule ordering constraints are necessary. In the constructional treatment of this problem, only two dimensions of structure are necessary — the word structure dimension and the phonetic structure dimension. We will illustrate the constructional approach by looking at the phonology of the Mohawk word meaning "I will push it," as indicated in Figure 3, where the W and P dimensions are given, W on top and P on bottom.

W: ye + ʌ̃k + hrek + ʔ #

P: y ɔ́ k r e g e ʔ #

Figure 3: W and P forms of "I will push it."

In the traditional generative treatment of this example, there would be a derivation beginning with what we have given as the W-form and ending with the P-form. Six rules would have to apply in sequence, three of them ordered in a particular way. The derivation would have seven lines. Figure 4 contains such a traditional generative derivation.

In Figure 4, the derivation is given on the left, and the rules are stated on the right. The second, third, and fourth rules must apply in that order. Here are the reasons: Stress assignment must precede epenthesis, since the stress rule in its most general form assigns stress to the second vowel from the end of the word. In this case, epenthesis (the third rule) inserts a vowel near the end of the word, and in this word the stress comes three vowels from the end. The general form of the stress rule can be preserved if it applies before epenthesis. Similarly epenthesis must precede the voicing of stops between vowels, since the epenthesis rule introduces one of the vowels (the /e/ following the /g/) that forms the environment for the voicing rule.

In generative phonology, the steps of a derivation are constructed one at a time, and these is no possibility of 'looking ahead' to a later stage of a derivation. But in constructional phonology, there are no intermediate stage, and no notion of before or after. There is simply a pairing of the W- and P-dimensions. Each construction plays the role of sanctioning a difference between the W- and P-dimensions.

In a construction, the environment conditioning the correspondence may appear in either the W- or P-dimensions. As we shall see, this obviates the need for rule ordering. Intuitively, a construction with an environment in the W-dimension will be translated into generative phonology as a relatively 'early' rule, while a construction with an environment in the P-dimension will be translated into generative phonology as a relatively 'late' rule. As we shall see, the Mohawk stress assignment construction has its environment (second vowel from the end of the word) stated in the W-dimension. When translated into generative phonology, it appears to apply 'before' epenthesis, which is a construction where the epenthetic vowel appears in the P-, but not the W-dimension.

Similarly, the environment in the epenthesis construction is in the W-dimension, while the environement for intervocalic voicing is in the P-dimension. Thus, when translated into generative phonology, epenthesis appears to apply 'before' voicing.

But in constructional phonology, there is no atemporal 'before' and 'after'. There are simply correlations across dimensions, in this case two dimensions. For the sake of exposition, we will build up the display of the correlations step-by-step by superimposing the constraints. Nothing will be changed. More constraints will be added. The ordering of exposition is purely arbitrary, and the superimposition should be thought of as holding 'all at once'.

Figure 5 contains the representations to be paired. Figure 6 contains a sequence of constraint applications in two parts: on the right is a construction; on the left, the constraint imposed by that construction is drawn as a bold line. Just to show that there is nothing holy about the order of presentation of the constructions, I will begin by presenting the three constructions that must be ordered in generative phonology, but I will scramble the order of presentation, with intervocalic voicing presented first, then stress assignment, then epenthesis. Other constructions will be presented after those.

It is an accident of history that phonology has been done in the mid-to-late twentieth century in terms of symbol-manipulation rules and abstract atemporal derivations. However, it appears that symbol-manipulation phonology and such atemporal derivations are inconsistent with the function of neural networks, and so it would appear, with the operation of the brain. The intermediate stages of such derivations have no cognitive reality. Constructional phonology appears to be better suited to a theory of mind in which the mind IS the physical brain. Generative phonology seems to require a separation of mind and brain so that the mind (in which generative rules operate) permits 'operations' outside of real time, which the brain, of course, does not.

Constructional and generative phonology have not yet been compared for their linguistic adequacy. Such comparisons should be forthcoming before long.

References

Brugman, Claudia. 1981. *Story of Over.* Master's Thesis, University of California, Berkeley.

Fauconnier, Gilles. 1985. *Mental Spaces.* Cambridge: MIT Press.

Lakoff, George. 1987. *Women, Fire, and Dangerous Things.* Chicago: University of Chicago Press.

Langacker, Ronald. 1987. *Foundations of Cognitive Grammar.* Stanford: Stanford University Press.

Regier, Terry. 1988. Recognizing Image-Schemas Using Programmable Networks. In this volume.

Ullman, Shimon. 1983. Visual Routines. A.I. Memo No. 723, Artificial Intelligence Laboratory, MIT.

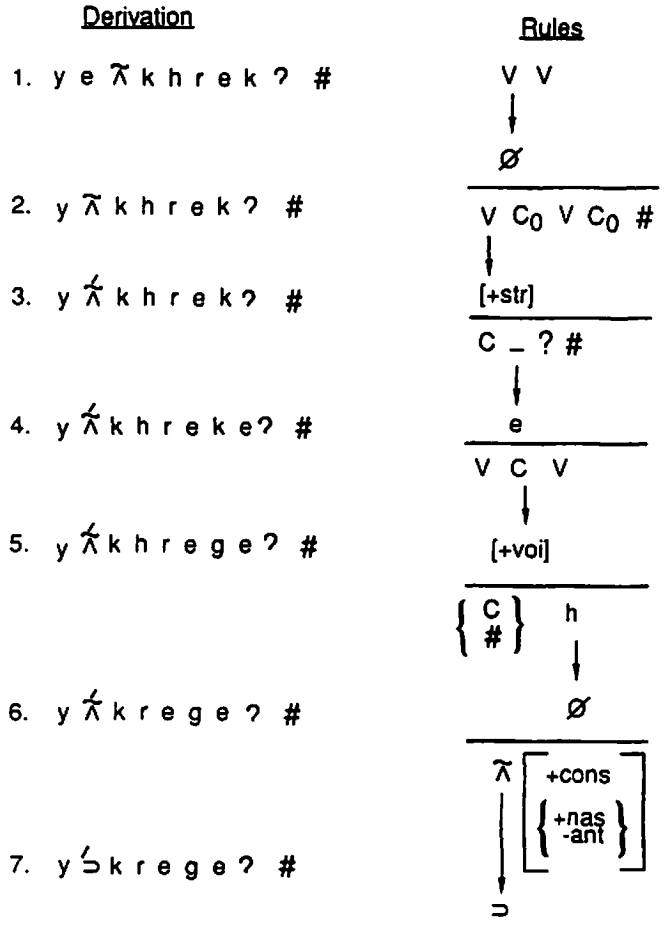

Figure 4: Traditional generative derivation of "I will push it."

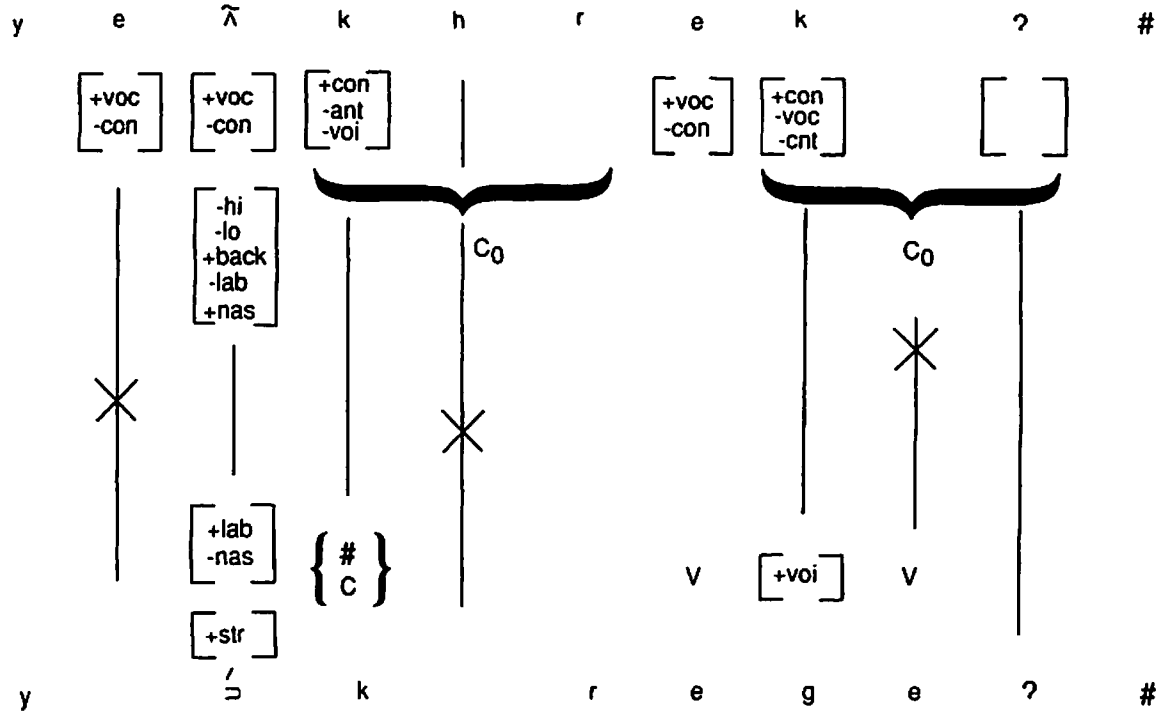

Figure 5: Pairing of word-level and phonetic forms.

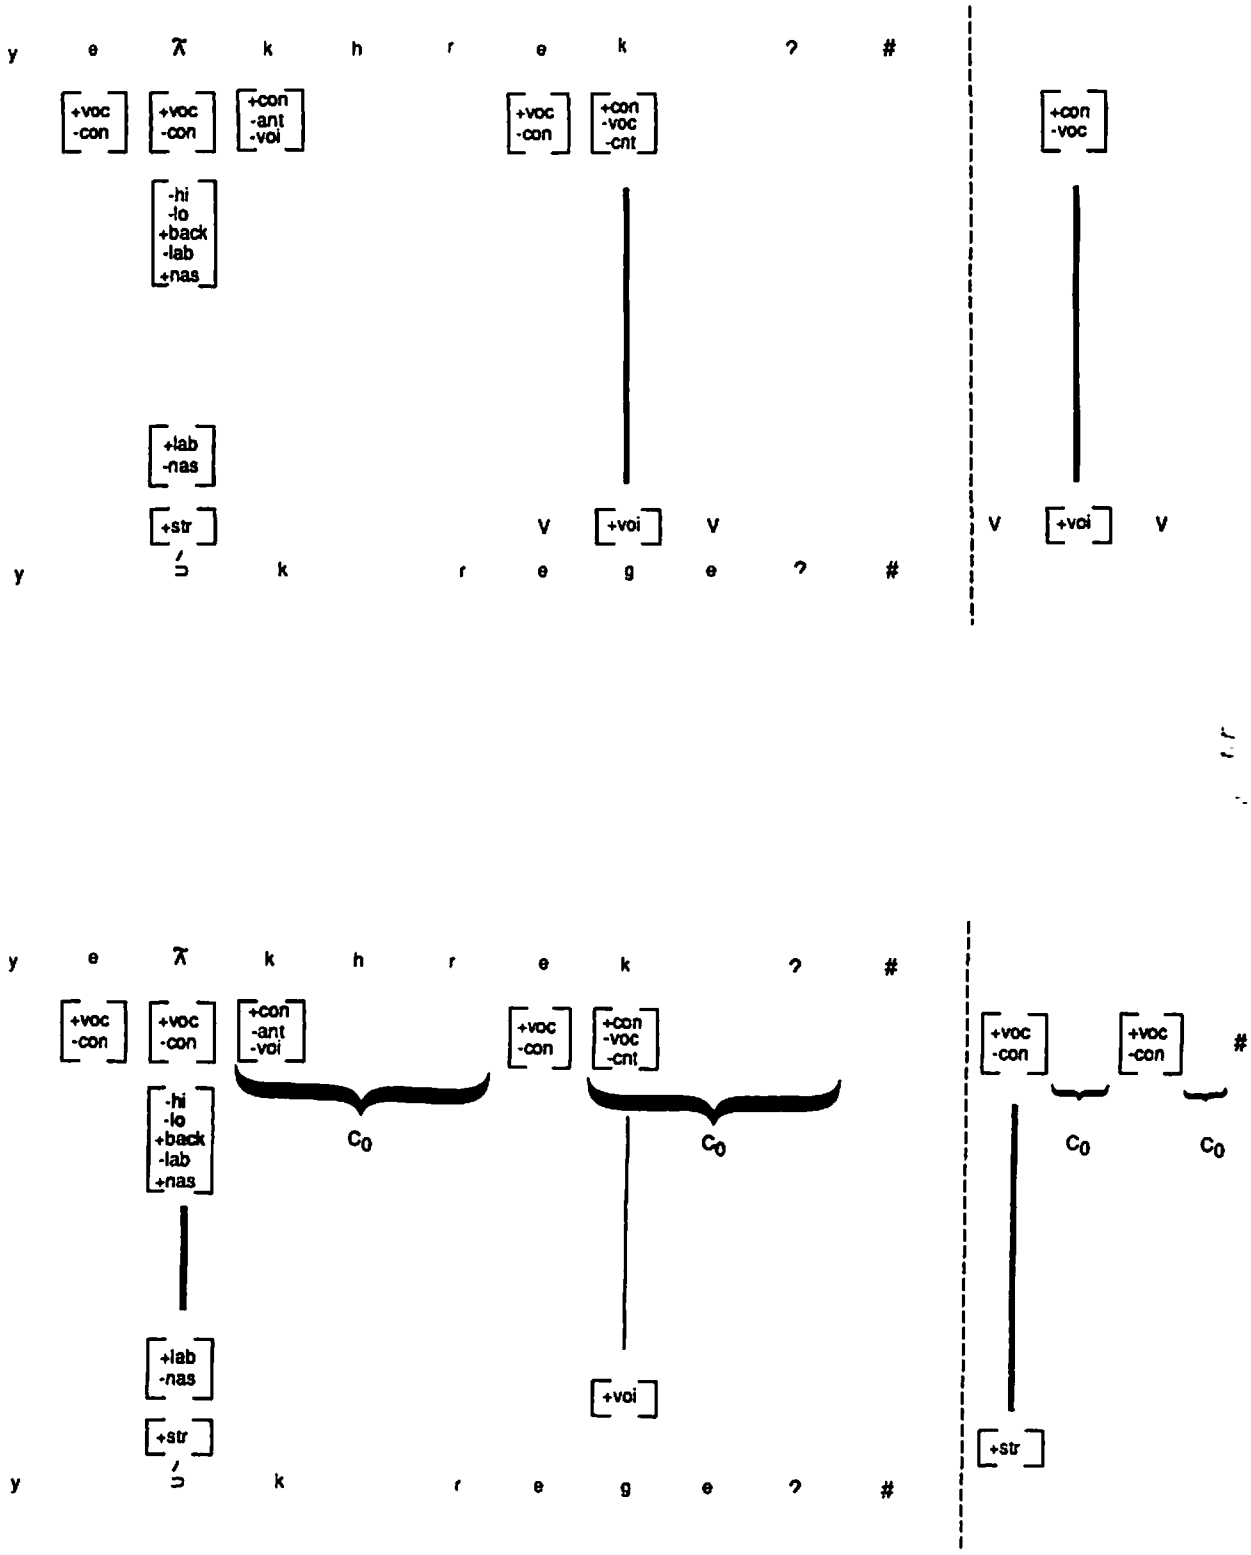

Figure 6: Constraints participating in the pairing of word-level and phonetic forms for "I will push it."

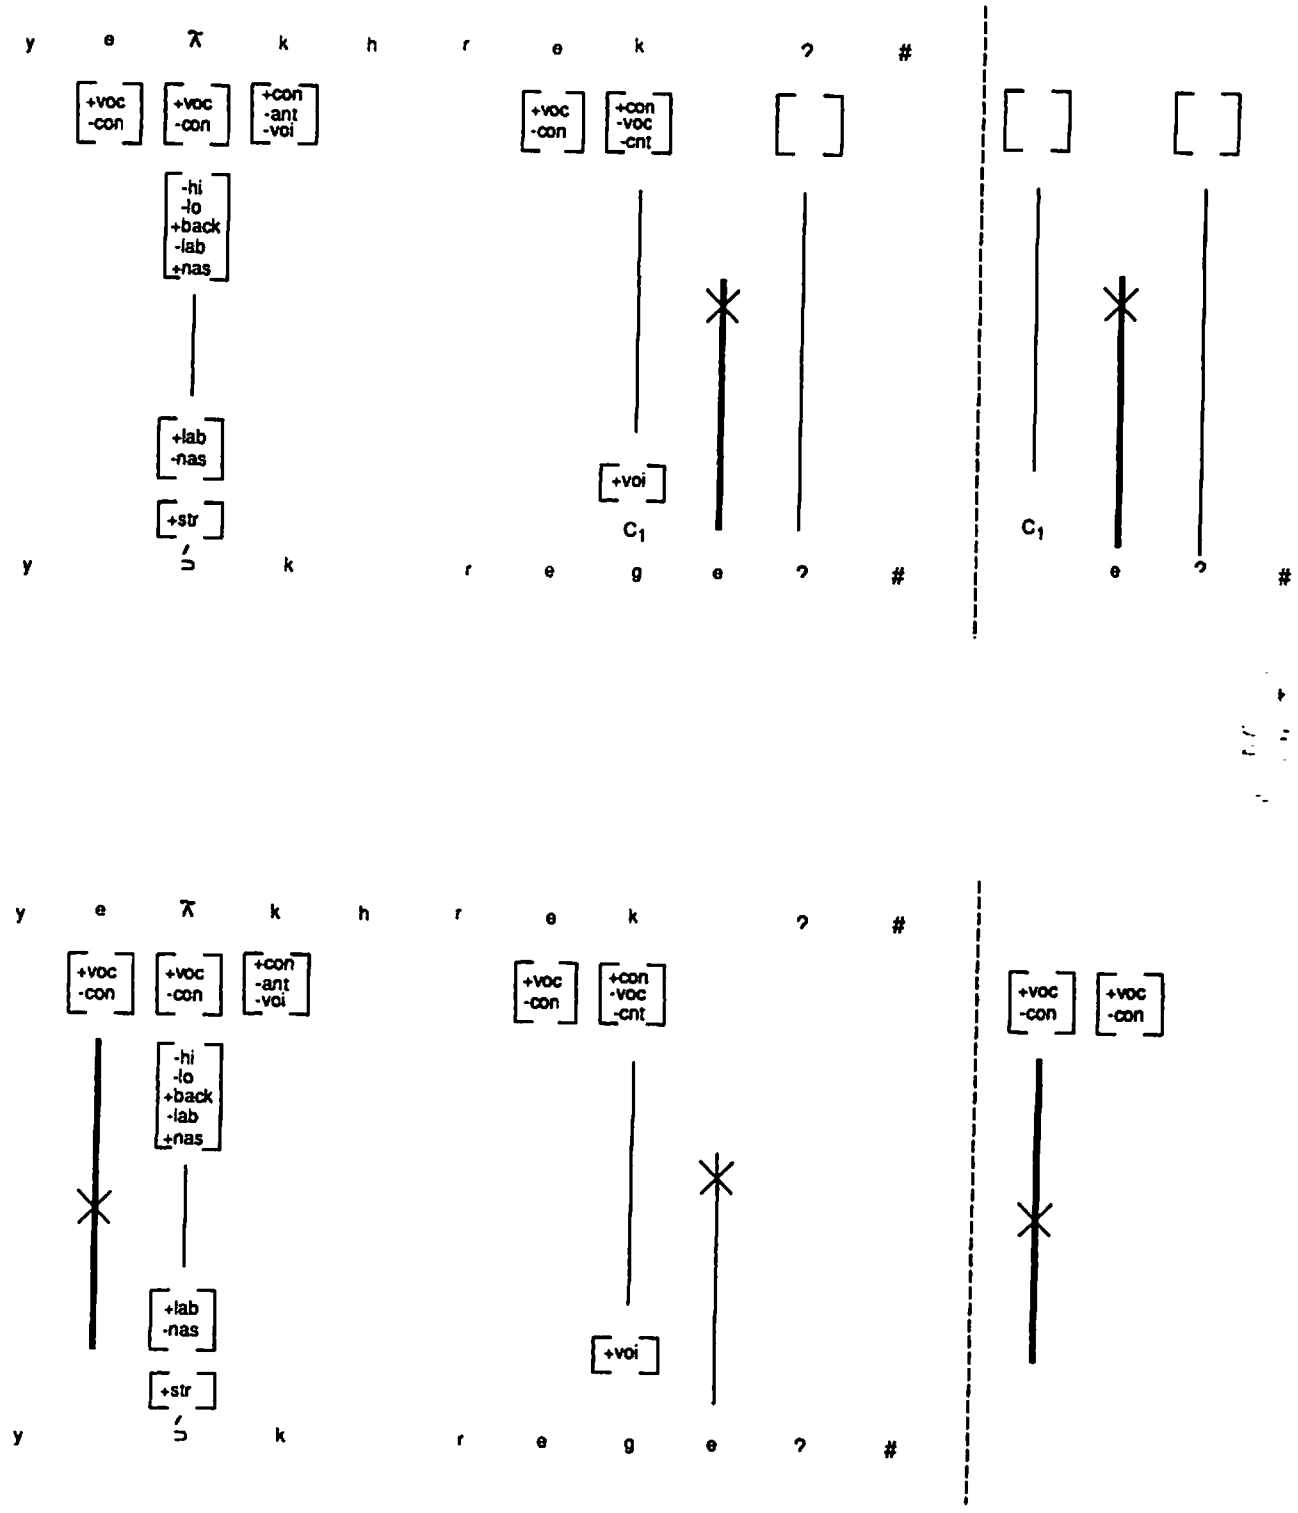

Figure 6 (continued)

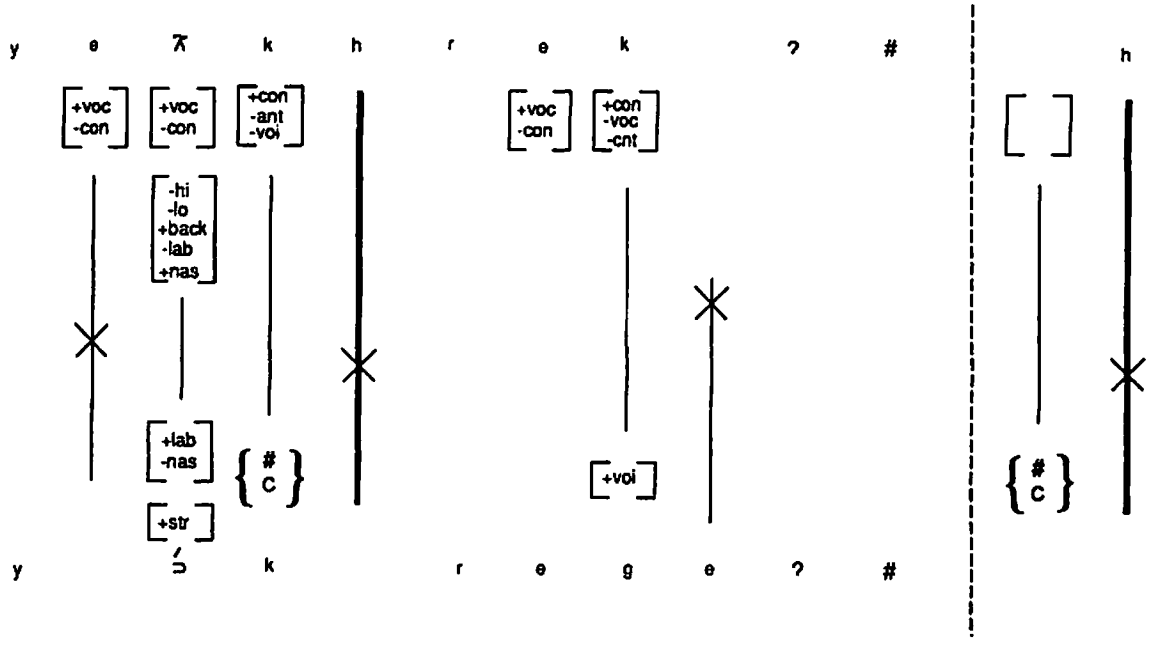

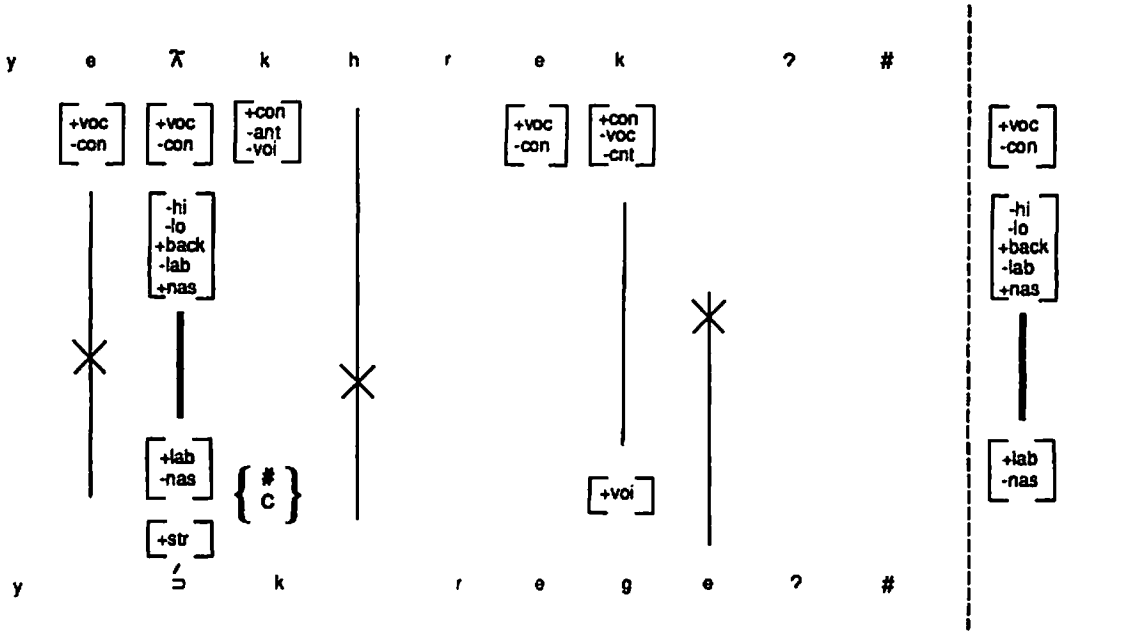

Figure 6 (continued)