Emergentist Approaches to Language

Brian MacWhinney
Carnegie Mellon University, Psychology

In Bybee, J. & Hopper, P. (Eds.) 2001. Frequency and the emergence of linguistic structure. Benjamins: New York.

It is easy to understand why many linguists are becoming attracted to the view of language as an emergent behavior. For over forty years, syntacticians have worked to establish a fixed set of rules that would specify all the grammatical sentences of the language and disallow all the ungrammatical sentences. Similarly, phonologists have been trying to formulate a fixed set of constraints that would permit the possible word formations of each human language and none of the impossible forms. However, neither language nor human behavior has cooperated with these attempts. Grammars keep on leaking, language keeps on changing, and humans keep on varying their behavior. Frustrated by these facts, linguists have begun to question the methodology that commits them to the task of stipulating a fixed set of rules or filters to match a specific set of data. Searching for more dynamic approaches, they have begun to think of language as an emergent behavior.

Some linguists worry that emergentism can distract us from the hard work of linguistic description. It would certainly be a mistake to abandon structured linguistic description without providing a solid mechanistic alternative. Emergentism is fully committed to providing empirically testable, mechanistic descriptions. However, discovering the exact shape of emergent mechanisms is no small task and it would be foolhardy to abandon traditional linguistic description before solid emergentist alternatives have been formulated. We need to understand what emergentism can offer us, while maintaining a certain skepticism regarding its immediate applicability. In order to begin to organize our thinking about emergent processes in language, the first question that we need to ask is "Emergence from what?" In other words, we need to be able to see how linguistic behavior in a target domain emerges from constraints derived from some related external domain. For example, an emergentist account may show how phonological structures emerge from physiological constraints on the vocal tract. This account invokes external determination, since the shape of one level of description is determined by patterns on a different level. Similarly, an emergentist syntactic account may show how variations in word order arise from patterns of morphological marking.

Emergence plays an important role in all of the physical and biological sciences. Consider the formation of the honeycomb. When a bee returns to the hive after collecting pollen, she deposits a drop of wax-coated honey. Initially, each of these honey balls is round and has approximately the same size. As these balls get packed together, they take on the familiar hexagonal shape that we see in the honeycomb. There is no gene in the bee that codes for hexagonality in the honeycomb, nor is there any overt communication regarding the shaping of the cells of the honeycomb. Rather, this form is an emergent consequence of the application of packing rules to a collection of honey balls of roughly the same size, as suggested in Figure 1.

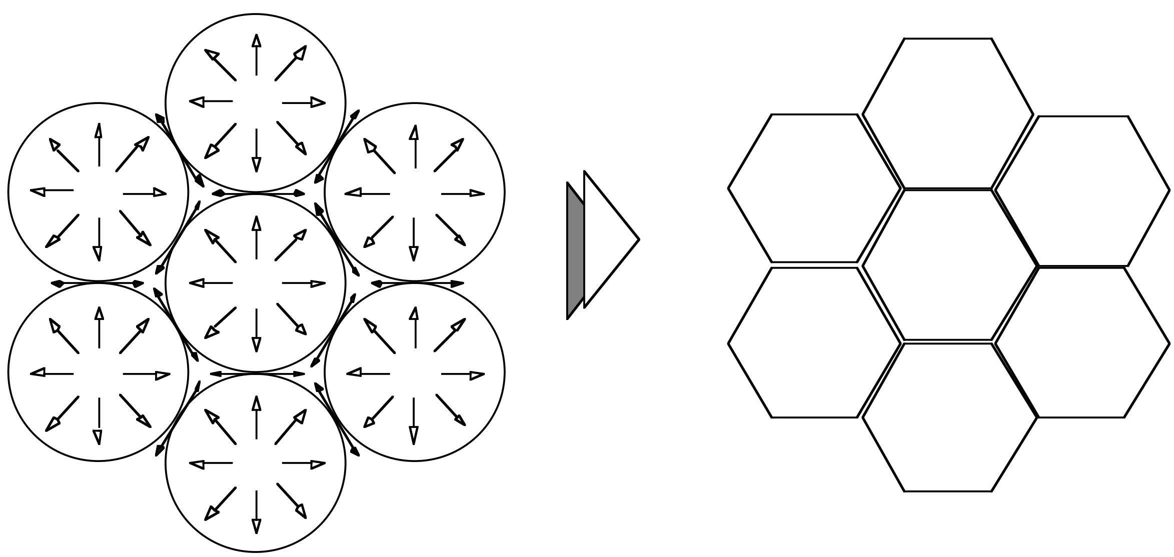

Figure 1: The emergence of hexagons in a honeycomb from the packing of spheres

Nature abounds with examples of emergence. The outlines of beaches emerge from interactions between geology and ocean currents. The shapes of crystals emerge from the ways in which atoms pack into sheets. Weather patterns like the Jet Stream or El Niño emerge from interactions between the rotation of the earth, solar radiation, and the shapes of the ocean bodies. Biological patterns emerge in much the same way. For example, the pattern of a leopard's spots is laid down in the first two days of embryonic development by the diffusion of two morphogens across the surface of the embryo. Variations in the patterns of stripes and dots on the skin emerge as consequences of the developing geometry of the embryo. Using a single-parameter reaction-diffusion physical model of a cylindrical embryo of varying sizes, Murray {, 1988 #9040} was able to simulate the emergence of marking patterns on the tails of the leopard, cheetah, jaguar, giraffe, zebra, and genet. The only parameter required for these simulations was the shape of the prenatal tail at 40 days. Similarly, Murray could model the shape of spots on the necks of different species of giraffe using what is known about variations in the shape of the embryo at 40 days.

Similar forces determine the emergence of patterns in the brain. For example, Miller, Keller, and Stryker {, 1989 #5066} have shown that the ocular dominance columns described by Hubel and Weisel {, 1963 #7114} in their Nobel-prize-winning work may emerge as a solution to the competition between projections from the different optic areas during synaptogenesis in striate cortex (see Figure 2).

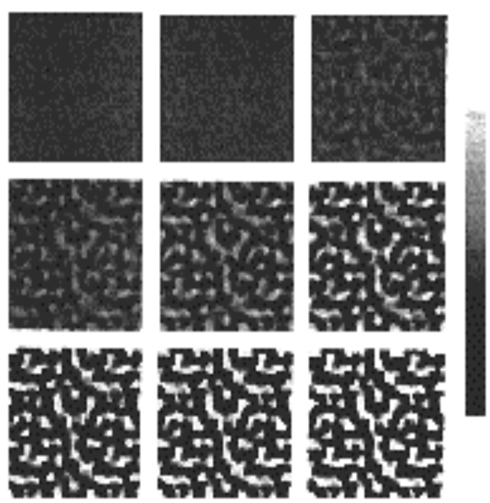

Figure 2: The emergence of ocular dominance columns, based on Miller et al. {, 1989 #5066}

Emergentist accounts of brain development provide useful ways of understanding the forces that lead to neuronal plasticity, as well as neuronal commitment. For example, Ramachandran {, 1995 #7421} has shown that many aspects of reorganization depend upon the elimination of redundant connectivity patterns. Moreover, Quartz and Sejnowski {, 1997

#7200} have shown that plasticity may also involve the growth of new patterns of connectivity. On the macro level, recent fMRI work {Booth, 1999 #8994} has shown how children with early brain lesions use a variety of alternative developmental pathways to preserve language functioning.

1. Levels of emergence

The emergentist accounts developed in the current symposium have focused on how frequency determines linguistic structure. In order to better understand the psychological bases of these analyses, we need to conduct a fundamental analysis of the types of emergent processes and the ways in which each are subject to the pressures of frequency, reliability, and other measures of cue validity. To begin this process of analysis, we can distinguish six separate temporal frames or levels for emergence.

1. **Evolutionary emergence**. The slowest moving emergent processes are those which are encoded in the genes. These processes, which are subject to more variability and competition than is frequently acknowledged, are the result of glacial changes resulting from the pressures of evolutionary biology. We can refer to this type of emergence as "evolutionary emergence". Language is a species-specific ability that depends, in part, on unique genetic patterns that have developed across the last five million years. However, it is unlikely that these emergent patterns directly code specific linguistic structures. Rather, all of these patterns have their effects filtered by the second level of emergence – epigenetic emergence.
2. **Epigenetic emergence**. Differential expression of embryonic DNA triggers a further set of processes from which the structure of the organism emerges {Gilbert, 1994 #9033}. Some physiological structures are tightly specified by particular genetic loci. For example, the recessive gene for phenylketonuria or PKU begins its expression prenatally by blocking the production of the enzymes that metabolize the amino acid phenylalanine. Although the effects of PKU occur postnatally, the determination of this metabolic defect emerges prenatally in terms of the production of particular enzymes. Other prenatal emergent anatomical structures involve a role for physical forces in the developing embryo. The formation of the spots on the leopard is an example of this type. Epigenetic effects continue after birth, as the processes of gene

expression interact with the ongoing physical and neurological changes in the organism. Some of these late-emerging processes may have important implications for the development of language. For example, the myelinization of neurons {Lecours, 1975 #2462} or the commitment of cerebral areas to stimulus processing {Blakemore, 1974 #9034; Julesz, 1995 #7413} are effects that arise epigenetically.

Emergentist accounts formulated on these first two scales are not fundamentally different from explanations that have figured in nativist theories. However, nativist theories have often failed to view these processes as emergent and have seldom distinguished between evolutionary and epigenetic emergence. By formulating nativist theory in emergentist terms, we gain a richer picture of the actual dynamic processes that shape human development. The next four levels of emergentist accounts also rely heavily on biology as the underpinning for self-organization. However, they allow for the unfolding of biological forces in more flexible and interactive fashions than those envisioned in the first two time scales.

3. **Emergence from local maps**. Accounts on this level emphasize the ways in which linguistic structures emerge from the local architectures of neural networks. We know that the cells of the cortex are organized into a series of columnar processing units including perhaps 100,000 cells in each unit. Within each processing unit, the organization of information obeys strict map-like patterns. Visual information is organized retinotopically, auditory information tonotopically, and motor information by individual limbs and digits. The formation of these local neural architectures is an emergent phenomenon, determined by processes such as inductance, the preference for short connections, cell differentiation, cell migration, competition for input, and lateral inhibition. Self-organizing feature maps (SOFM) provide a particularly useful way of expressing our current knowledge of this local level of neural structure. Many of the properties of human language emerge from the ways in which input is processed by local feature maps. Clear examples of this type of emergence include the Pierre-Humbert model of phonetic entrenchment (this volume), the Bybee model of morphological entrenchment (this volume), or the various connectionist models of

the acquisition of morphology. Models on this level deal with issues such as chunks, dual-processing, gang effects, and exemplar-based processing.

4. **Emergence from functional circuits**. High-level cognition arises from the interaction of local processing units across long distances in the brain. Cortical processing in local maps is gated and amplified by signals from the thalamus, hypothalamus, hippocampus, amygdala, cerebellum, and basal ganglia. Within the cortex, frontal areas such as the cingulate, the dorso-lateral prefrontal cortex, and Broca's area work to modify the processing of posterior language areas in the temporal and parietal lobes. As patterns become transmitted across longer distances in the brain, temporal constraints start to place limits on information storage and retrieval. In order to deal with these limitations, systems such as the phonological loop {Gathercole, 1993 #6961} or the output monitor {Shattuck-Hufnagel, 1979 #3763} use functional neural circuits to maximize performance. Properties of these functional circuits determine many aspects of the shape of human language, particularly on the levels of syntax and discourse. Examples of models based on the operation of these circuits include Baddeley's {, 1992 #5837} articulatory loop, the Carpenter and Just CC-CAPS model of language processing {, 1992 #5180}, Anderson's rational model of cognition {, 1993 #5762}, or the Competition Model {MacWhinney, 1989 #5822}.
5. **Grounded emergence**. Although models based on local maps and functional circuits are well-grounded in neuronal terms, they cannot express the ways in which language functions in a real social context {Vygotsky, 1962 #4273; Goffman, 1974 #1563}. Nor can they capture effects that are determined by the fact that the speaker has a real body {MacWhinney, 1999 #7785}. The groundings provided by the social context and the body provide two further sources for the emergence of language structure. Social forces and the shape of the ongoing conversation embed language in a framework of givenness, topicality, backgrounding, coreference, and shared knowledge that facilitates successful communication {Givón, 1979 #1533}. Accounts that explore these forces include conversation analysis, discourse analysis, and much of sociolinguistics. At the same time, we use the projection of our own perspectives onto the experiences around us to extract personalized meaning from

social interactions {MacWhinney, 1999 #7785}. By taking and shifting perspectives, we can assimilate objects, space, time, causation, and social frames to our own physicalist mental models. Accounts that explore these forces include Cognitive Grammar {Bailey, 1997 #8089} and various new developments in psychology that could be called Embodiment Theory.

6. **Diachronic emergence**. The changes that languages undergo across centuries can also be viewed in emergentist terms. Some diachronic processes tend to level distinctions and contrasts, others introduce new forms and contrasts {Bybee, 1988 #608}. Just as erosion and orogeny work together to determine the geologic landscape, forces of leveling and innovation work together to determine the changing linguistic landscape. Among the most important processes are regularization {Bybee, 1985 #607}, entrenchment {Brooks, 1999 #9039}, gang effects {Hare, 1995 #7172}, lexical innovation {Clark, 1979 #823}, semantic bleaching, and phonological neutralization (Pierrehumbert, this volume).

This paper will focus on these last four types of emergence. These are the levels of emergence that have figured most prominently in recent psycholinguistic research and modeling.

2. Emergence from Local Maps

Connectionist models use nodes, connections, and activation to model the processing of information in local networks. These models come in many types, including Boltzmann machines, back propagation nets, recurrent nets, Hopfield nets, and Kohonen nets {Fausett, 1994 #6891}. Although the bulk of work in the modeling of language processes has used back propagation nets, there are some known limitations to this particular architecture {Grossberg, 1987 #5522}. An interesting alternative to back propagation is the Kohonen network or self-organizing feature map (SOFM) {Miikkulainen, 1993 #6971}.

The most important feature of the self-organizing feature map is its ability to encode lexical items in an emergentist, but still localist fashion. Although the position of a lexical item in a field is determined by a distributed pattern of features in a sparse matrix, these features still reliably activate a consistent node or area of nodes in the map. Figure 3 shows how the semantic fields for a few common nouns become self-organized. In this figure, we see that words that share semantic features are close to each other in the semantic map. For example, the verb *hit* is close to *broke* and the noun *lion* is close to *dog*. On the phonological or lexical map, monosyllables are grouped together on the right and disyllables on the left. This patterning is a consequence of the phonological coding chosen for this particular simulation. If another system of phonological features has been used, a different pattern of similarity would have emerged. The important point is that proximity of any two items on the map is determined by the similarity of their featural representations.

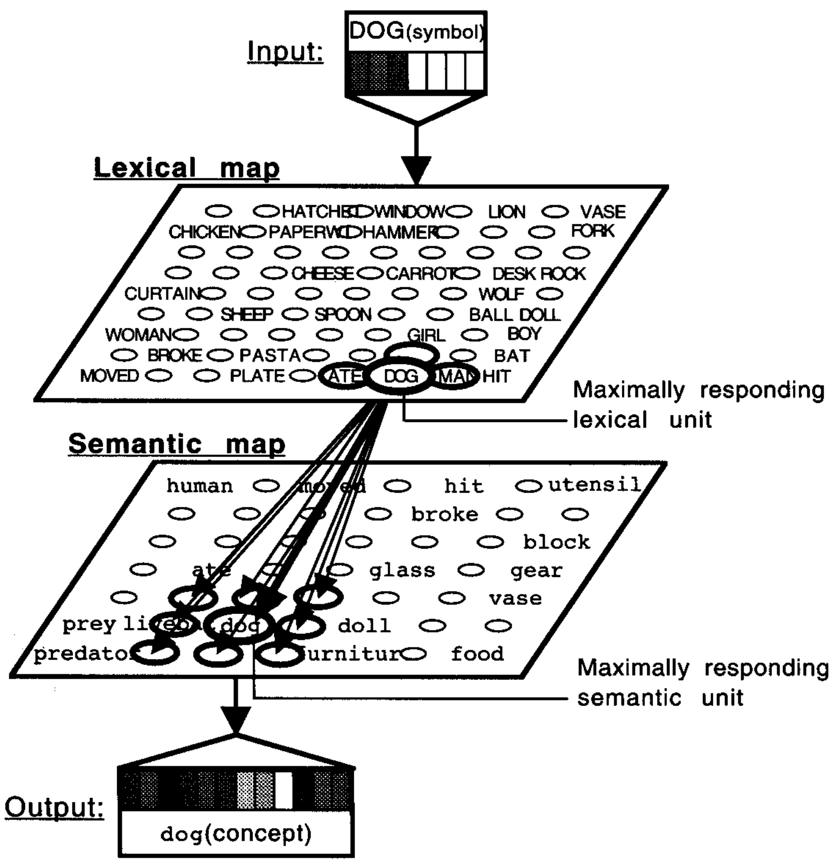

The lexicon architecture. The lexical input word DOG is translated into the semantic representation of dog. The representations are vectors of gray-scale values between 0.0 and 1.0, stored in the weights of the units. The size of the unit on the map indicates how strongly it responds. Only a few strongest associative connections of the lexical unit DOG are shown.

Figure 3: From Miikulainen {, 1993 #6971}, this map illustrates the emergent activation of the phonological form of the word *dog* on the lexical map and the meaning of *dog* on the semantic map.

Miikkulainen {, 1993 #6971} has shown how a wide range of linguistic phenomena, from polysemy to the parsing of relative clauses, can be explained within the framework of the self-organizing feature map. Feature maps rely on a system of lateral inhibition between nodes that closely mimics actual biological processes found in many areas of the cortex. Moreover, these networks can also be constructed in a way that emphasizes the brain's preference for the maintenance of short connections. Extending Miikkulainen's work, Li and MacWhinney {, 1999 #8645} have shown how these maps can learn the meaning and semantic applicability of the reversive prefixes in English to produce correct forms such as *disassemble* or *unbutton* as well as overgeneralizations such as *unappear* or *disfasten*. The input to this simulation used semantic feature codes derived both from rating studies with subjects and vectors from the HAL (Hyperspace Analogue to Language) database of Burgess and Lund {, 1997 #7853}. HAL represents word meanings through multiple lexical co-occurrence constraints in large text corpora. Words are coded using a string of 100 numbers in which each number represents a value on a statistically-extracted semantic dimension.

Feature maps provide a method for encoding the emergence of individual lexical items. In back propagation models, it is impossible to identify a structure that corresponds to a lexical item. This is because lexical items are represented by a distributed pattern of features. Feature maps also use distributed representations as input. However, because they emphasize the emergence of a topology of similarity, specific lexical items develop a clear identity. At first, a word may match a fairly large area in feature map space, such as an area with a six-unit radius. However, as the learning of additional words progresses, the radius devoted to that item decreases. Toward the end of learning, words come to compete specifically with their neighbors and it is this competition that sharpens the topological separation between lexical items. The emergence of a linkage between lexical items and a position on a map does not involve any overt "writing" of lexical labels on localist nodes {Stemberger, 1985 #3987; Dell, 1986 #1029}. Instead, the association of an item to an area in the map is an emergent process. In fact, some items move around a bit on the map during the first stages of learning.

Feature maps can control the three basic linguistic processes of rote, combination, and analogy. The Dialectic Model {MacWhinney, 1978 #2690} recognized these three processes as central to accounts of language acquisition. However, the formulation of a neural network model that deals with each of these three processes has proven difficult. First let us consider how feature maps deal with the process of rote learning.

Unlike many other neural network systems, feature maps are capable of "one-shot" associative learning. This means that they can learn a new word on a single trial without unlearning earlier forms. Feature maps share their ability to handle one-shot learning with a few other neural network architectures, such as SDM {Kanerva, 1988 #6942} and ART {Grossberg, 1987 #5522}. The ability to handle one-shot learning is crucial, because it permits exemplar-based learning. Exemplar-based learning models are superior in various ways to those that do not make a clear encoding of examples {Corrigan, 1988 #922; Tomasello, 1992 #6719; Goldberg, 1999 #8629}. For example, Kruschke's {, 1992 #5463} ALCOVE model of concept learning is grounded on the learning of examples. Taraban {, 1993 #5504} has shown how an exemplar-based model is needed to capture the earliest stages of the learning of Russian gender marking or the learning of new forms in a Miniature Linguistic System. Similarly, Matessa and Anderson {Matessa, 2000 #8987} have compared ACT-R and the Competition Model. They show that, in miniature linguistic system experiments by McDonald and MacWhinney {, 1991 #2870}, as well as in a new experiment designed specifically to compare the two models, ACT-R does a better job of predicting the order of cue acquisition. The reason for the better performance of ACT-R is that it focuses learning on one cue at a time, whereas the Competition Model processes all cues at all times during learning. This cue focusing allows ACT-R to quickly acquire frequent cues and to initially block learning about less frequent cues. In this way, ACT-R does a better job of modeling actual human learning.

The ability to model one-shot learning allows a network to model much of what we have begun to learn about the role of frequency in promoting rote, chunking, and entrenchment. As Bybee, Corbett et al. (this volume), Frisch (this volume), Hare (this volume), MacWhinney, Marchman, Pierrehumbert, Plunkett, and many others have argued, high frequency allows forms to become entrenched. However, as Corbett et al. (this volume) and Frisch (this volume) have shown, neural network models must assign correct values to the contrasting effects of token frequency, type frequency, construction frequency, and paradigm frequency. In order to model frequency effects on each of these levels, our models have to provide a role for each of these levels of structure. However, these levels themselves should be viewed as emergent. For example, the development of a unique phonology for phrasal chunks such as *I don't know* {Bybee, 1999 #9095} underscores the importance of mechanisms for acquiring frequent phrasal units.

The second major process invoked by the Dialectic Model {MacWhinney, 1978 #2690} is analogy. Because of the distributed nature of their input representations, feature maps do a good job of modeling analogic processes. Because neighborhood structure is based on featural similarity, feature maps can model the various prototype effects and gang effects that are usually captured by neural network models.

The third major process invoked by the Dialectic Model {MacWhinney, 1978 #2690} is combination. One of the simplest types of combination is the attachment of a suffix to a stem to mark a category such as plural or past in English. In recent years, Pinker {, 1991 #6945}, Clahsen {, 1999 #8816}, Marslen-Wilson {Marslen-Wilson, 1998 #8660}, and others have underscored the importance of default patterns in morphology. Attempts to model even this basic level of combination in neural networks have met with mixed results. The problem is that the formulation of a model that includes rote, analogy, and combination in a single architecture requires more complexity than can be found on a local map. We will discuss ways of constructing such an architecture when we examine the joining of local maps into functional neural circuits.

Before leaving the topic of local maps, it is important to mention the potential role for neuronal recruitment and reorganization in emergentist models. Following a suggestion of Miikkulainen {, 1993 #6971}, Ping Li and I have been exploring an extension of feature maps based on the notion of map sprouting as a result of competition. The idea is as follows. As the child learns more and more words, the principal lexical feature map starts to become overcrowded. To deal with this competition, words that are close competitors project their competition to a secondary neural area which is designed specifically to handle competitions between smaller sets of words. For example, the cohort of words beginning in /kæ/ could project to a single area. These would include *cat*, *catalog*, *catastrophe*, *cab*, *California*, *candle* and *cattle*. Although these words would still have a representation on the main feature map, the importance of that representation would diminish over time as the secondary map took over the competition. All that the main map would continue to process would be the basic onset syllable structure or BOSS {Taft, 1981 #4064}. This same type of recruitment of secondary arenas for competition can occur on both the semantic and phonological level, as illustrated in Figure 4. A mechanism of this sort can help us understand how phonological and semantic categories emerge during the normal course of word learning.

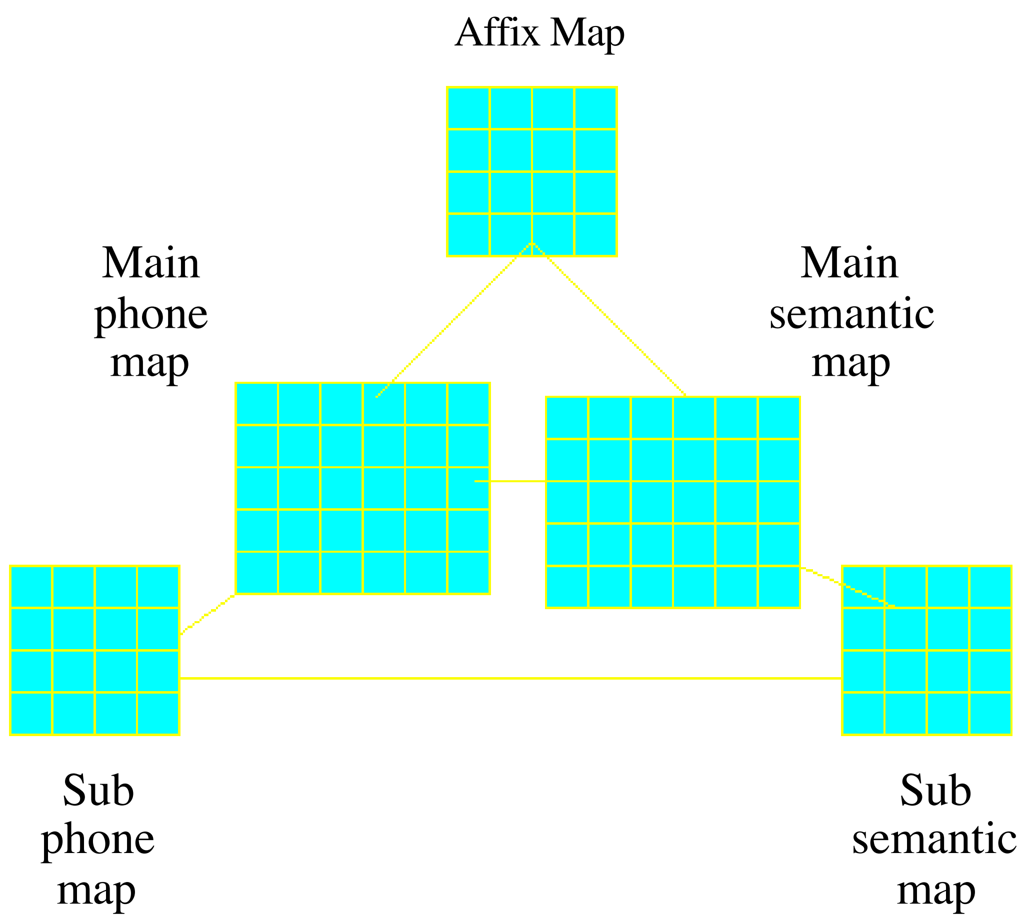

Figure 4: The emergence of secondary processing areas to resolve cohort competition

3. Chunking

Neural network models make no claims regarding the shape of phonological and semantic inputs. They assume that the shape of these inputs is determined by perceptual mechanisms that lie outside of the scope of the core simulation. However, changes in the shape of the input can radically alter the outcome of learning in neural networks. One aspect of input representations that needs to be carefully explored is the extent to which speakers process words in terms of phrasal chunks, rather than more analytic morphemes. The tendency of both children and adults to process high frequency phrases as units has been discussed in terms of the process of chunking by researchers such as Bybee, Bush, Boyland, and Scheibman (this volume). Although it is clear that chunking plays a major role in language learning and processing, it is important to clarify several issues that arise in these discussions.

1. The term "chunk" can refer to unitization in perception, production, or memory. In models such as ACT-R {Anderson, 1993 #5762} or SOAR {Newell, 1990 #5300}, chunks are the basic units of declarative encodings. However, these models make clear internal distinctions between chunks in perception, production, and memory. When we are operating outside of the explicit framework of these models, it is probably confusing to use a single term for all three levels of unitization. Instead, we can consider using terms such as "Gestalt" or "perceptual chunk" for units in perception and "avalanche" or "motoric chunk" for units in production. The term "Gestalt" is tightly linked to perceptual processes. The term "avalanche" {Grossberg, 1978 #6512; Gupta, 1997 #6908} refers to a series of units that have been chained together for output production. Avalanches are serial strings of behaviors in which the triggering of the beginning of the string leads to the firing of all its component pieces. Thus, the avalanche is used to control production of words or even phrases.
2. We may believe that chunks arise both through perceptual chunking and avalanche formation. One fact that argues for this analysis is the observation that the exact shape of reductions is often highly lexically specific. For example, in the phrase *I don't know*, the deletion of the first flap is specific to this particular phrase. Similarly, the reduction of *What's up with you*? to / relies heavily on a precise mapping to the original phrasal form. One way of explaining this assumes that reductions first arise through simplificatory processes in production, but are then stored by perceptual processes that are unique to the phrasal item. The crucial assumption here is that feature maps can use whole perceptual chunks as their inputs. This form of processing would be used to account not only for phrases such as *I don't*

know but also for common nominal phrases or constructions of the type that show lexical effects for French liaison (Bybee, this volume). Neural networks have not yet been used to model these effects.

3. The reductions that occur in avalanches can have negative perceptual consequences. For example, Vroomen and de Gelder {, 1999 #8985} have shown that phoneme monitoring for initial segments is more difficult in words that have been resyllabified in fluent speech. Given this, listeners must develop ways of dealing with the problems caused by chunking effects in production. The problem is that many phrases appear in both a fluent unitized form and a more analytic, less chunked form. This means that the perceptual system needs to be able to recognize both forms when required. Recognition of unitized forms is facilitated by the fact that they are typically high in frequency.

4. Emergence from Functional Circuits

The consolidation of information in chunks in local maps is an important component of language learning and processing. However, no small set of local maps can process the rich complexity that is contained in even the simplest sentences. In order to develop more complex neural circuitry, the brain must have ways of connecting local maps into larger functional circuits. Hebbian learning provides one way of establishing such connections. For Hebbian learning to work properly between local maps, it is necessary that the maps be as least partially interconnected. We can refer to these interconnections between local maps as long distance connections. In Hebbian learning, long distance connections will be strengthened when the units to which they are connected fire at the same time. This means that connections between nodes that do not fire together will weaken and disappear over time. This type of learning works well for the formation of links between feature maps. For example, the /kaet/ node in the phonological map will tend to fire at the same time as the *cat* node in the semantic map. This will lead to the strengthening of the connection between the two nodes on the two maps. The presence of the connection is a given, but its relative strength is emergent. Moreover, there is reason to believe that the connection itself could emerge when needed {Quartz, 1997 #7200}. This type of long distance mapping probably involves connections between temporal auditory areas and temporo-parietal semantic areas. When the child comes to linking up words to potential articulations, even more distant connections must be established to frontal areas in motor cortex and Broca's area for speech planning.

4.1. Three models

One example of a model that deals with the formation of these connections between areas is the Gupta and MacWhinney {, 1997 #6908} model of the development of articulatory forms in the child. This model links together the concept of an articulatory plan or "avalanche" {Grossberg, 1978 #6512} with the notion of a feature map. The architecture of the model is given in Figure 5.

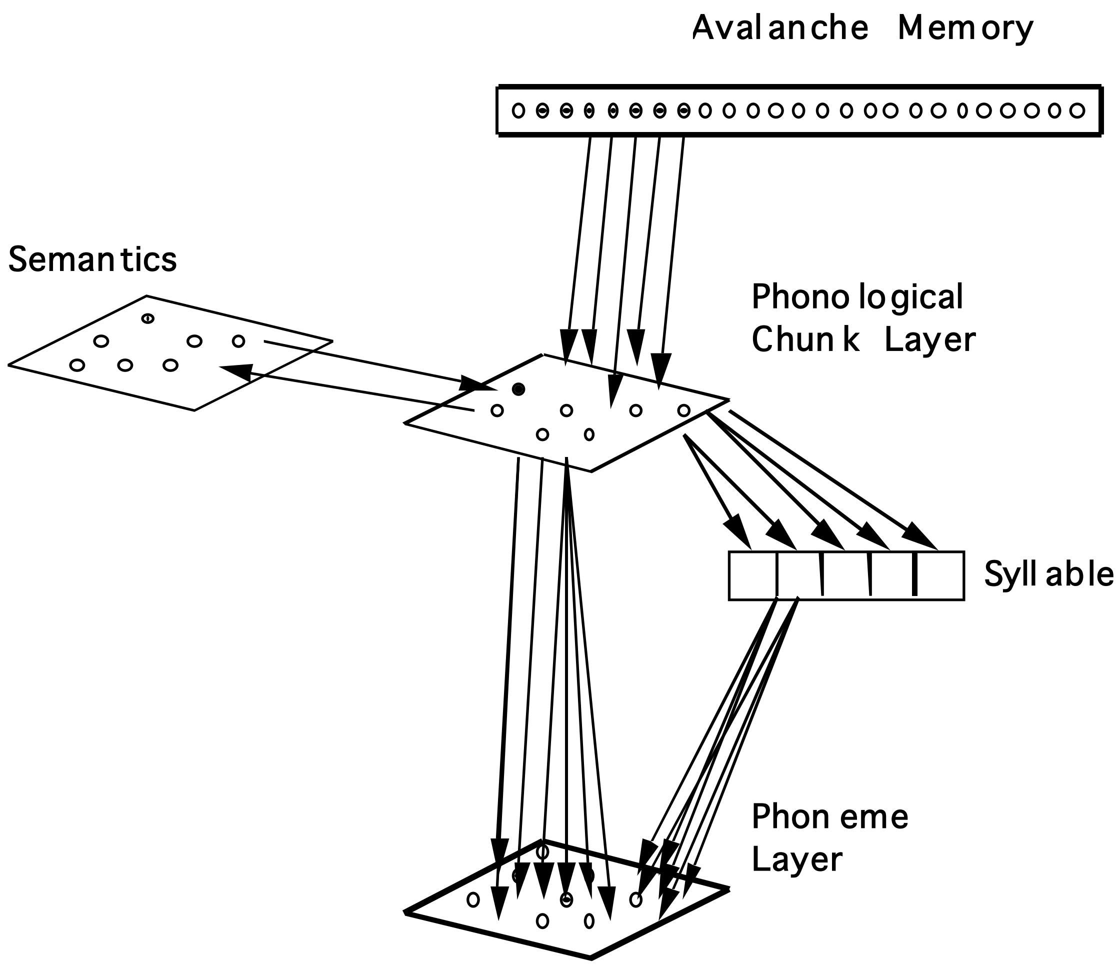

Figure 5: The model of Gupta and MacWhinney (1999) for learning of articulatory forms

In this figure, words are represented as stored strings or avalanches. The phonological chunk layer is a feature map with pointers to each individual avalanche. It also maintains connections to the phoneme layer that facilitates the recognition of syllabic templates. As in the model of Figure 3, a layer of semantic connections organizes phonological processing.

A model developed by Plaut and Kello {, 1999 #8634} provides another example of a how language form emerges from connections between processing areas. This model shows how articulatory form emerges from attempts to match input phonology during babbling and the learning of the first words. In this system, a series of six connections between processing areas are used to allow the sounds of words to train the formation of articulations.

A third model {MacWhinney, 1999 #7833} explains how syntactic processing can be derived from more distant connections between local feature maps. That model uses a core structure in which the semantic and phonological maps of Figure 3 are dependent on a third map of central lexical forms. From these central lexical forms, there are then connections not only to the semantic map, but also to an output phonology map (as in Figure 5) and an input phonology map. In addition, lexical items have connections to phrases or constructions in another map. This model is not yet implemented.

All three of these models link local processing fields into larger functional circuits. As they stand, all three models are preliminary and incomplete. However, they illustrate how complex functional circuits can be built up using local maps as their components.

4.2. Processing effects

Current models of sentence processing focus on the ways in which lexically-based constructions provide cues for role assignments. The assignment of sentence elements to particular grammatical roles is performed through a competitive process based on the relative strength of the cues involved {MacWhinney, 1989 #5822}. The Competition Model uses various measures of cue reliability to predict cue strength in experiments in which cues are placed in competition. The notion of reliability developed in this work is essentially the conditional probability of an interpretation, given a cue. If the interpretation is always correct when the cue is present, this probability approaches 1.0. For example, in the Italian sentence, "Il spaghetti mangia Giovanni" (The spaghetti eats Giovanni), the noun *spaghetti* competes with the noun *Giovanni* for the role of subject of the verb *mangia*. The cue that favors *spaghetti* is its initial positioning in the NVN order, whereas the cue that favors *Giovanni* is its animacy. In Italian, animacy is a stronger and more reliable cue than word order and so the sentence is given an OVS interpretation. In English, the opposite is true, since word order is more reliable than animacy. Thus, in English, we end up with an implausible interpretation of an event in which some animated spaghetti wants to eat Giovanni.

The basic result of Competition Model work has been that the most reliable cues in a language are also the strongest ones in sentence processing. The relative dominance order of cues varies markedly across languages and is closely tuned to reliability. In addition, cue strengths function additivel, so that an array of interacting weak cues can sometimes dominate over one cue with medium validity. However, no combination of weak cues can ever dominate over a truly strong and reliable cue. These patterns have been observed in dozens of studies in children, adults, aphasics, and bilinguals speaking 15 different languages. The view of sentence processing as dependent on cue validity has since been widely supported by other recent work in psycholinguistics {Trueswell, 1994 #7220; Tanenhaus, 1989 #7389; MacDonald, 1994 #7187; MacDonald, 1999 #8628}.

Recent psycholinguistic work has supported the probabilistic and competitive assumptions of the Competition Model; it has also underscored the extent to which syntactic competition emerges directly from individual lexical constructions. For example, MacDonald et al. {, 1994 #7187} show how a lexically-based version of the Competition Model can be used to account for the processing of lexical ambiguities, including prepositional phrase attachment, main verb vs. reduced relative competitions, and direct object vs. complement clause ambiguity. Consider the processing of the ambiguity in the garden-path sentence "The horse raced past the barn fell." Initially, *raced* is interpreted as a main verb in the past tense. However, the suffix *–ed* has a secondary reading as a marker of the past participle. When the verb *fell* is encountered, the interpretation of *raced* as the verb of the main clause encounters competition. To resolve this competition, the past perfect reading of *–ed* is strengthened and a reduced relative interpretation is constructed.

Although reliability is an excellent predictor of eventual sentence interpretation, we now know that the actual on-line processing of syntactic cues is also strongly influenced by the forces of frequency or availability. Listeners come to rely initially on cues that are always present, even if they are not uniformly reliable. For example, in Russian, listeners are willing to wait for the eventual case cue, since it will be reliable when it is encountered {Kempe, 1999 #8082}. In German, on the other hand, listeners just decide to go with what they have, since no single cue is all that reliable or universally available.

4.3. Frequency Effects

The contrast between reliability and frequency effects discussed in the previous section underscores the importance of paying careful attention to the exact shape of frequency effects in sentence processing and language change. Although frequency effects are pervasive throughout language, the targets of these effects need to be carefully specified. Consider these issues:

1. It is generally accepted that a form becomes stronger when it occurs more frequently. However, for this to work, the system has to detect new instances of a form as related to old instances. This means that the system must perform a similarity match. If a new input closely resembles a previous input, it will activate as the winner its closest match in the map. If the input lies between two currently strong nodes, the system has to be tuned to allow it to emerge as a new center of activation or new lexical item. These effects work in a similar fashion on both segmental and lexical levels. Thus, categorization emerges as a property of the design of neural networks and the way that they process frequency information. This issue arises particularly when the system is attempting to deal with phrasal simplifications such as *supchu* or the reduced form of *I don't know*. If it attempts to map these items onto their component pieces, it may end up misperceiving in other less idiosyncratic cases.
2. Should our counting of frequency apply to tokens, types, or collocations? Within the context of feature map theory, both types and tokens must be counted. Tokens have their effect through repeated activation of the same type nodes. Types have their effects through neighborhood effects. For example, a given conjugational pattern may be frequent in terms of the types of verbs to which it applies, but not particularly frequent in terms of the actual number of tokens to which it applies. This will occur when the pattern applies to a large number of fairly infrequent stem types. Most neural network models have not yet dealt with frequency effects that are due to

constructions. In order to capture such effects, it will be necessary to elaborate the view of these models in terms of functional neural circuits, as discussed earlier.

3. What is the effect of frequency on pattern productivity? The debate about the status of default inflections as "rules" {Bybee, 1995 #5892} may reduce to a discussion of the technical parameters that need to be set in a neural network model to model productivity for patterns with a high type applicability.
4. To what extent can frequency preserve old structures? On the one hand, old structures are preserved against leveling by frequency. On the other hand, the fact that these resistant forms are no longer in accord with new patterns tends to open them to semantic reinterpretation, as in the development of *went* as the past tense of *go*.
5. What is the effect of transition probability on fusion, contraction, and affixation? The merger of highly frequent combinations in production leads, over time, to their reinterpretation and acquisition as single forms over time.
6. What is the effect of frequency on sound change? Sound change has typically been viewed as operating across the board. Flege (in press) has recently shown that sound changes in second language learning also work in this way. However, Phillips (this volume) has shown that sound change affects high frequency items first. What are the mechanisms driving this relation?
7. What is the effect of frequency on semantic bleaching or other functional changes? According to the Competition Model {MacWhinney, 1989 #2725; MacWhinney, 1997 #7449}, each grammatical device is itself a coalition of functional motives or pressures that exist in a peaceful coexistence. Although the subject of an English sentence might express definiteness 75% of the time, it might also express perspective 95% of the time. However, if other forces start to tip this balance, we could see a progressive association of subjecthood with definiteness. Over time, subject marking could be identified not as a way of coding perspective, but as a way of coding definiteness. Other examples of reinterpretation include the fusion of *what*, *is*, and *up* to form *sup*. In these cases, as it becomes impossible to extract the original morphemes, the meaning of the merged unit starts to shift. Forms like *goodbye* or

even *zounds* represent the end result of this process of reinterpretation of merged forms.

5. Grounding

Local neural maps can account for many fundamental effects in language usage. If we supplement these local mechanisms with functional neural circuits, we can account for still more aspects of parsing, syntax, and language production. Although this neural circuitry provides many of the mechanisms that support cognition and language, a full account must go beyond neurons and circuits. Much of the actual content of cognition is grounded in our bodies and our social lives. Meaning arises from the fact that our minds are embedded in our bodies that experience motion, vision, hearing, and emotions through our sensory organs and muscles. At the same time, we act as social agents who are embedded in ongoing conversations that determine and facilitate the shape of cognition.

MacWhinney {, 1999 #7785} examines the issue of symbol grounding by linking linguistic form to perspective-taking. According to this analysis, when we listen to sentences, we engage in an active process of role-taking by assuming the perspective of the grammatical subject. From this perspective, we begin to interpret the actions, objects, and positions involved in the sentence. Grammatical devices such as relativization, passivization, topicalization, pronominalization, and switch-reference all serve to direct the process of perspective-taking through various perspective shifts. On the lowest level, these processes involved deictic {Ballard, 1997 #7835} identifications of objects in memory. We process these objects in terms of their physical affordances {Gibson, 1977 #7939}. We use perspective-switching to coordinate multiple perspectives and frames in space and time that are marked through aspectual and spatial language. Perspective also allows us to interpret the causal actions involved in transitive constructions {Hopper, 1980 #1938}.

Social perspective-taking allows us to shift between competing social frames {Fauconnier, 1996 #7559}. In both narrative and conversation, we attempt to coordinate a wide array of referents into a set of coherent perspectives. We then shift back and forth between these perspectives in order to construct social reality. These effects are illustrated in Thompson and Hopper's account (this volume) of the actual usage of transitive markings in conversation, as well as Sheibman's examination (this volume) of perspectival effects on person-marking in conversation.

Functional accounts of perspective-shifting have a variety of antecedents {MacWhinney, 1977 #2689; Firbas, 1964 #1300; Langacker, 1995 #7927; Chafe, 1974 #693}. However, recent advances in cognitive neuroscience {Rizzolatti, 1996 #7776; Kosslyn, 1995 #7568} are now showing us exactly how perspective-taking is implemented in the brain. As our understanding of these mechanisms grows, we will develop a clearer idea of how language emerges from physical and social perspective-taking.

6. Summary

Our tour of the different levels of emergentist accounts has helped us examine three basic issues:

1. **Emergence from what**? We have seen that the use of emergentist theories depends very heavily on the temporal level of the processing involved. Some accounts refer to child language development; others refer to language processing; yet other refer to language change. For each of these types of emergence, very different forces are at work.
2. **Frequency of what**? We have seen that neural networks are able to encode a wide variety of frequency effects. Some of these effects apply to articulations; others apply to lexical items; yet others apply to constructions. These effects include chunking in production, reinterpretation, overgeneralization, and resistance to overgeneralization.
3. **Integration**. Our models of language usage need to integrate levels, although many phenomena can be addressed on a single level. Integrated models will need to link frequency effects to the deeper processes of grounding in social relations, perspective-taking, consciousness, and the movements of the human body.

The articulation of emergentist accounts provides us with exciting new ways of linking linguistic theory to the rest of the human sciences.

References

Anderson, J. (1993). *Rules of the mind*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Baddeley, A. (1992). Working memory: The interface between memory and cognition. *Journal of Cognitive Neuroscience, 4*, 281-288.

Bailey, D., Feldman, J., Narayanan, S., & Lakoff, G. (1997). Modeling embodied lexical development. *Proceedings of the 19th Meeting of the Cognitive Science Society*, 18-22.

Ballard, D. H., Hayhoe, M. M., Pook, P. K., & Rao, R. P. (1997). Deictic codes for the embodiment of cognition. *Behavioral and Brain Sciences, 20*, 723-767.

Blakemore, C., & van Sluyters, R. (1974). Reversal of the physiological effects of monocular deprivation in kittens: Further evidence of a sensitive period. *Journal of Physiology, 237*, 195-216.

Booth, J. R., MacWhinney, B., Thulborn, K. R., Sacco, K., Voyvodic, J., & Feldman, H. (1999). Functional organization of activation patterns in children: Whole brain fMRI imaging during three different cognitive tasks. *Progress in Neuropsychopharmocology and Biological Psychiatry, 23*, 669-682.

Brooks, P. J., Tomasello, M., Dodson, K., & Lewis, L. B. (1999). Young children's overgeneralizations with fixed transitivity verbs. *Child Development, 70*, 1325-1337.

Burgess, C., & Lund, K. (1997). Modelling parsing constraints with high-dimension context space. *Language and Cognitive Processes, 12*, 177-210.

Bybee, J. (1985). *Morphology: A study of the relation between meaning and form*. Amsterdam: John Benjamins.

Bybee, J. (1995). Regular morphology and the lexicon. *Language and Cognitive Processes, 10*, 425-455.

Bybee, J., & Scheibman, J. (1999). The effect of usage on degrees of constituency: The reduction of *don't* in English. *Linguistics, 37*, 575-596.

Bybee, J. L. (1988). Semantic substance vs. contrast in the development of grammatical meaning. *Berkeley Linguistics Society, 14*.

Chafe, W. (1974). Language and consciousness. *Language, 50*, 111-132.

Clahsen, H. (1999). Lexical entries and rules of language: A multidisciplinary study of German inflection. *Behavioral and Brain Sciences, 22*.

Clark, E. V., & Clark, H. H. (1979). When nouns surface as verbs. *Language, 55*, 767-811.

Corrigan, R. (1988). Who dun it? The influence of actor-patient animacy and type of verb in the making of causal attributions. *Journal of Memory and Language, 27*, 447-465.

Dell, G. (1986). A spreading-activation theory of retrieval in sentence production. *Psychological Review, 93*, 283-321.

Fauconnier, G., & Turner, M. (1996). Blending as a central process of grammar. In A. Goldberg (Ed.), *Conceptual structure, discourse, and language* (pp. 113-130). Stanford, CA: CSLI.

Fausett, L. (1994). *Fundamentals of neural networks*. Englewood Cliffs, NJ: Prentice Hall.

Firbas, J. (1964). On defining the theme in functional sentence. *Travaux Linguistiques de Prague, 1*, 267-280.

Gathercole, V., & Baddeley, A. (1993). *Working memory and language*. Hillsdale, NJ: Lawrence Erlbaum Associates.

Gibson, J. J. (1977). The theory of affordances. In R. E. Shaw & J. Bransford (Eds.), *Perceiving, acting, and knowing: Toward an ecological psychology* (pp. 67-82). Hillsdale, NJ: Lawrence Erlbaum.

Gilbert, S. F. (1994). *Developmental Biology. Fourth edition*. Sunderland, MA: Sinauer.

Goffman, E. (1974). *Frame analysis*. NewYork: Harper and Row.

Goldberg, A. E. (1999). The emergence of the semantics of argument structure constructions. In B. MacWhinney (Ed.), *The emergence of language* (pp. 197-213). Mahwah, NJ: Lawrence Erlbaum Associates.

Grossberg, S. (1978). A theory of human memory: Self-organization and performance of sensory-motor codes, maps, and plans. *Progress in Theoretical Biology, 5*, 233-374.

Grossberg, S. (1987). Competitive learning: From interactive activation to adaptive resonance. *Cognitive Science, 11*, 23-63.

Gupta, P., & MacWhinney, B. (1997). Vocabulary acquisition and verbal short-term memory: Computational and neural bases. *Brain and Language, 59*, 267-333.

Hare, M., & Elman, J. L. (1995). Learning and morphological change. *Cognition, 56*, 61-98.

Hopper, P. J., & Thompson, S. A. (1980). Transitivity in grammar and discourse. *Language, 56*, 251-299.

Hubel, D., & Weisel, T. (1963). Receptive fields of cells in striate cortex of very young, visually inexperienced kittens. *Journal of Neurophysiology, 26*, 994-1002.

Julész, B., & Kovacs, I. (Eds.). (1995). *Maturational windows and adult cortical plasticity*. New York: Addison-Wesley.

Just, M., & Carpenter, P. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review, 99*, 122-149.

Kanerva, P. (1988). *Sparse distributed memory*. Cambridge, MA: MIT Press.

Kempe, V., & MacWhinney, B. (1999). Processing of morphological and semantic cues in Russian and German. *Language and Cognitive Processes, 14*, 129-171.

Kosslyn, S. M., Thompson, W. L., Kim, I. J., & Alpert, N. M. (1995). Topographical representations of mental images in primary visual cortex. *Nature, 378*, 496-498.

Kruschke, J. (1992). ALCOVE: an exemplar-based connectionist model of category learning. *Psychological Review, 99*, 22-44.

Langacker, R. (1995). Viewing in grammar and cognition. In P. W. Davis (Ed.), *Alternative linguistics: Descriptive and theoretical models* (pp. 153-212). Amsterdam: John Benjamins.

Lecours, A. R. (1975). Myelogenetic correlates of the development of speech and language. In E. H. Lenneberg & E. Lenneberg (Eds.), *Foundations of language development: A multidisciplinary approach* (Vol. 1, pp. 121-136). New York: Academic Press.

Li, P. (1999). Generalization, representation, and recovery in a self-organizing neural network of language acquisition. In M. Hahn & S. C. Stoness (Eds.), *Proceedings of the 21st Annual Meeting of the Cognitive Science Society* (pp. 308-313). Mahwah, NJ: Lawrence Erlbaum Associates.

MacDonald, M. (1999). Distributional information in language comprehension, production, and acquisition: Three puzzles and a moral. In B. MacWhinney (Ed.), *The emergence of language* (pp. 177-196). Mahwah, NJ: Lawrence Erlbaum Associates.

MacDonald, M. C., Pearlmutter, N. J., & Seidenberg, M. S. (1994). Lexical nature of syntactic ambiguity resolution. *Psychological Review, 101*(4), 676-703.

MacWhinney, B. (1977). Starting points. *Language, 53*, 152-168.

MacWhinney, B. (1978). The acquisition of morphophonology. *Monographs of the Society for Research in Child Development, 43*, Whole no. 1, pp. 1-123.

MacWhinney, B. (1989). Competition and lexical categorization. In R. Corrigan, F. Eckman & M. Noonan (Eds.), *Linguistic categorization* (pp. 195-242). Philadelphia: Benjamins.

MacWhinney, B. (1999a). The emergence of language from embodiment. In B. MacWhinney (Ed.), *The emergence of language* (pp. 213-256). Mahwah, NJ: Lawrence Erlbaum.

MacWhinney, B. (Ed.). (1999b). *The emergence of language*. Mahwah, NJ: Lawrence Erlbaum Associates.

MacWhinney, B., & Bates, E. (Eds.). (1989). *The crosslinguistic study of sentence processing*. New York: Cambridge University Press.

MacWhinney, B., & Pléh, C. (1997). Double agreement: Role identification in Hungarian. *Language and Cognitive Processes, 12*, 67-102.

Marslen-Wilson, W., & Tyler, L. K. (1998). Rules, representations, and the English past tense. *Trends in Cognitive Sciences, 2*, 428-435.

McDonald, J. L., & MacWhinney, B. (1991). Levels of learning: A microdevelopmental study of concept formation. *Journal of Memory and Language, 30*, 407-430.

Miikkulainen, R. (1993). *Subsymbolic natural language processing*. Cambridge, MA: MIT Press.

Miller, K., Keller, J., & Stryker, M. (1989). Ocular dominance column development: Analysis and simulation. *Science, 245*, 605-615.

Murray, J. D. (1988). How the leopard gets its spots. *Scientific American, 258*, 80-87.

Newell, A. (1990). *A unified theory of cognition*. Cambridge, MA.: Harvard University Press.

Pinker, S. (1991). Rules of Language. *Science, 253*, 530-535.

Plaut, D. C., & Kello, C. T. (1999). The emergence of phonology from the interplay of speech conrephension and production: A distributed connectionist approach. In B. MacWhinney (Ed.), *The emergence of language* (pp. 381-416). Mahwah, NJ: Lawrence Erlbaum Associates.

Quartz, S. R., & Sejnowksi, T. J. (1997). The neural basis of cognitive development: A constructivist manifesto. *Behavioral and Brain Sciences, 20*, 537-596.

Ramachandran, V. S. (1995). Plasticity in the adult human brain: Is there reason for optimism? In B. Julesz & I. Kovacs (Eds.), *Maturational windows and adult cortical plasticity* (pp. 179-198). New York: Addison-Wesley.

Rizzolatti, G., Fadiga, L., Gallese, V., & Fogassi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research, 3*, 131-141.

Shattuck-Hufnagel, S. (1979). Speech errors as evidence for a serial-ordering mechanism in sentence production. In W. E. Cooper & E. C. T. Walker (Eds.), *Sentence processing: Psycholinguistic studies presented to Merrill Garrett* (pp. 295-342). Hillsdale, N. J.: Lawrence Erlbaum.

Stemberger, J. (1985). *The lexicon in a model of language production*. New York: Garland.

Taft, M. (1981). Prefix stripping revisited. *Journal of Verbal Learning and Verbal Behavior, 20*, 289-297.

Tanenhaus, M., Carlson, G., & Trueswell, J. C. (1989). The role of thematic structures in interpretation and parsing. In G. T. M. Altmann (Ed.), *Parsing and interpretation* (pp. SI 211-234). Hove: Lawrence Erlbaum Associates.

Taraban, R., & Palacios, J. M. (1993). Exemplar models and weighted cue models in category learning. In G. Nakamura, R. Taraban & D. Medin (Eds.), *Categorization by humans and machines*. San Diego: Acdemic Press.

Tomasello, M. (1992). *First verbs: A case study of early grammatical development*. Cambridge: Cambridge University Press.

Trueswell, J. C., & Tanenhaus, M. K. (1994). Toward a lexicalist framework for constraint-based syntactic-ambiguity resolution. In J. C. Trueswell & M. K. Tanenhaus (Eds.), *Perspectives in sentence processing* (pp. 155-179). Hillsdale, NJ: Lawrence Erlbaum Associates.

Vroomen, J., & de Gelder, B. (1999). Lexical access of resyllabified word: Evidence from phoneme monitoring. *Memory and Cognition, 27*, 413-421.

Vygotsky, L. (1962). *Thought and language*. Cambridge: MIT Press.