

# Exemplar models, evolution and language change<sup>1</sup>

ANDREW B. WEDEL

## *Abstract*

*Evidence supporting a rich memory for associations suggests that people can store perceptual details in the form of exemplars. The resulting particulate model of category contents allows the application of evolution theory in modeling category change, because variation in categorized percepts is reflected in the distribution of exemplars in a category. Within a production-perception feedback loop, variation within an exemplar-based category provides a reserve of variants that can serve as the seeds for shifts in the system over time through random or selection-driven asymmetries in production and perception. Here, three potential pathways for evolutionary change are identified in linguistic categories: pruning of lines of inheritance, blending inheritance and natural selection. Simulations of each of these pathways are shown within a simple exemplar-based model of category production and perception, showing how consideration of evolutionary processes may contribute to our understanding of linguistic category change over time.*

## **1. Introduction**

Although we all learn about Darwinian evolution in the context of change within biological species, the mechanism of evolution itself is independent of biology. Darwinian evolution is a population-based mechanism of change and

---

1. I would like to thank Samira Farwaneh, Joseph Salmons, Kenny Smith, Adam Ussishkin and two anonymous reviewers for valuable comments and suggestions. All errors remain mine.

can operate in any system that can be described as a reproducing population exhibiting a distribution of traits, where variants show different rates of reproduction in correlation with some trait, and where that trait is inherited at above-chance levels by offspring. Over time, differential reproduction rates of variants with different heritable traits enrich the population with traits that are associated with greater rates of successful reproduction. Because of the generality of these criteria, the last few decades have seen a growing recognition that the mechanisms of Darwinian evolution potentially operate in many systems outside of biological species (see, e.g., Dawkins 1982, 1983; Dennett 1995; Henrich and Boyd 2002; Henrich et al. in press). Not surprisingly, many linguists have noted as well that language as a system seems to fit these general criteria (e.g., Lindblom 1992; Ohala 1989, 1992; Labov 1994; Czikó 1995; Hurford 1999; Bybee 2001; Blevins 2004). First, language is not learned as an undifferentiated whole, but rather in bits and pieces over time. Furthermore, idiolectal variation and error ensure that any one of these bits of language, such as the pronunciation of a particular word, exists in the learner's data set as a distribution. As a consequence, the data from which a learner abstracts a language system can be understood as a population of variants. Given that internal and external factors conspire to make some variants more likely to be learned and imitated than others, and that persistence of learned patterns is usually long relative to generation time, the population of variants exhibited in a generation should be influenced by the pattern of variants learned by previous generations. In this article, I want to show you that exemplar models of language provide a particularly convenient tool for modeling language change via evolutionary processes, and in doing so, illustrate several potentially productive avenues for thinking about language change through evolution. Before beginning, however, it may be useful to revisit some of the arguments that evolutionary theory can be extended beyond its original purview.

The operation of Darwinian evolution in biological systems is so well-known that it can be difficult to mentally disentangle the properties of the infrastructure in biology that supports evolution from the necessary characteristics required for evolutionary change itself. For example, while many have argued that cultural practices can be subject to natural selection (see, e.g., Henrich and Boyd 2002), others have argued in return that because mental representations are neither as discrete nor as faithfully transmitted as genes, natural selection cannot be understood as more than a metaphor for change in the cognitive/cultural sphere (e.g., Lewontin 2005). However, as Henrich et al. (in press) argue, this all-or-nothing reasoning about evolution is likely to be misleading, because many of the conditions that contribute to evolution as a mechanism of change are themselves gradient, allowing evolutionary processes to contribute to change to a greater or lesser degree depending on the details of the system itself.

For example, one factor that influences the efficiency of natural selection is the degree to which the system approximates a population of discrete ‘replicators’ (see, e.g., Dawkins 1982). The ideal system that is usually considered when thinking about natural selection consists of fully discrete elements, each of which replicates fully independently of all other system elements, producing faithful copies of itself. In this case, selective pressure promoting replication of elements with a given characteristic has a maximally efficient influence on the composition of the subsequent generation. At the other extreme, if no part of a system can replicate without the rest, then selection cannot operate at all. Crucially however, intermediate positions along this scale also allow selection to function. Genes, for example, do not replicate fully independently of one another because they are physically linked together in larger units, such as chromosomes and organisms. Because of this, reproduction of a detrimental gene sequence can be linked to reproduction of an advantageous one, reducing the efficiency of selection of the advantageous gene sequence.

Further along the continuum away from full discreteness, we can imagine a system that has no real independent replicators, but simply consists of a distribution that replicates each point along the distribution in each generation by interpolating gradually between all nearby points. In this case, reproductive ‘units’ have no discernable boundaries, nor is any ‘offspring’ point the faithful copy of any ‘parent’ point. However, if some selection process can act to amplify or diminish the replicative contribution of some portions of the distribution over others, then this selection will result in a new distribution that has evolved in qualitatively the same sort of way that we would expect for a distribution of fully discrete replicators. A simulation of such a system and discussion of this point can be found in Henrich and Boyd (2002), and we will see a demonstration of selection in a system without fully discrete replicators below in Section 3.3.

A further example of a gradient condition on evolution is the degree of correlation between a unit of reproduction and a selectable characteristic. In the absence of any correlation at all, selection-driven change is impossible, but on the other hand, absolute correlation is not required either. We see this in biological evolution, where a gene can increase or decrease in a population even when it produces a selected trait only in rare circumstances. In general, the stronger the association between a gene and a trait, the stronger the selection can be, but because there are many degrees of association between genes and traits, the relative contribution of natural selection to the distribution of genes in a population can vary widely.

We can see then that a wide variety of systems can support evolutionary processes to some degree, and even the infrastructure of biological systems itself provides the conditions for evolution in partial and varying, rather than absolute ways. As a consequence, the observation that the infrastructure un-

derlying language is very unlike that of genetic systems is not sufficient for concluding that evolutionary theory cannot be applied to the problem of language change. Words are not the same as genes, and people as language users are not equivalent to people as members of a biological species, but the basic theory of Darwinian evolution is independent of the particularities of the biological systems in which it was first developed. Once we decide that language may supply some degree of the necessary conditions for evolution, a more useful, empirical question is whether evolutionary processes do in fact contribute in any significant ways to language change, and if so, how.

Despite the fact that linguistic and biological systems are different, linguists have not needed to start entirely from scratch in exploring how evolutionary processes may impact language change. Evolutionary biologists have been thinking for over a century about the pathways by which biological species can change through evolution, and the study of language change can benefit from their findings. In particular, when we identify abstract pathways of change that may be shared between linguistic and biological systems, we can generate hypotheses by examining the outcome of the biological pathway. Conversely, if we identify pathways for change in language that do not exist in biological systems, we know that the potential resulting patterns of change may not resemble patterns familiar from biology. In this article, I will focus on three pathways of change that may operate in language, two of which have clear correlates in biological evolution, and one which does not, and show how they can be illuminated within an exemplar model of language. I will expressly not try to argue that these mechanisms are necessarily operative in significant ways in linguistic evolution, but by showing that they are *potentially* operative under current models of language processing, I hope to highlight possible questions and lines of research.

Many modern treatments of language change (e.g., Hare and Elman 1995; Kirby and Hurford 2002; Blevins 2004) focus on the role of intergenerational transmission in the generation of new grammar variants. If these variants acquired by naïve learners within a speech community survive as differences in adult speech patterns, they influence the range and distribution of data presented to the next generation of learners, generating positive feedback which can result in a future, community-wide shift. However, since at least Paul (1886: Chapters 2 and 3) and Baudouin de Courtenay (1895/1972), linguists have recognized that adult speakers continue to be able to produce and recognize a wide variety of instances of the 'same' linguistic utterance, and that this provides for the possibility that the speech of an adult may itself change over time. In support of these early models of language change, a good deal of recent evidence shows that the fine details of linguistic categories can continue to evolve in response to experience after the initial stage of category formation in juvenile acquisition (e.g., Sancier and Fowler 1997; Goldinger

2000; Harrington et al. 2000; Norris et al. 2003; reviewed in Labov 1994; Bybee 2001; Pierrehumbert 2003). The mechanism(s) that underlie such gradient category change over the life of an individual are worth studying, because surface pattern shifts within adults may provide much of the variability that underlies larger scale phase transitions in language transmission (Cooper 1999).

Exemplar models provide a succinct theory of the life of a category within an individual adult speaker which allows us to conveniently model feedback-driven change in that category in response to experience (reviewed in Pierrehumbert 2002, 2003). This article argues that exemplar models of linguistic category structure are useful for this purpose, because they provide both an explicit account of category contents in terms of experience, and an easily modeled mechanism for individual experiences to influence future production and categorization behavior. I begin below with a summary of exemplar models and their properties as relevant to this article, followed by examples from simulations that illustrate several plausible pathways of change within adult speakers. Because of the present focus on exemplars, I limit discussion here to possible pathways of language change within individuals, although abstractly parallel pathways can frequently be identified at the level of transmission between generations.

## **2. Exemplar models**

The assumption of a fundamental distinction between general, abstract knowledge and specific, episodic memory has a long tradition in the psychological literature on categorization. In recent decades however, research has consistently shown that people detectably retain access to highly detailed memories of an event for a long time, and make use of these memories when carrying out categorization tasks previously thought to require only general knowledge (for reviews, see Tenpenny 1995; Johnson 1997). To account for these findings, a class of newer theories has developed which locates recognition and storage of specific, episodic memories at the core of categorization processes (Hintzman 1986; reviewed in Jacoby and Brooks 1984). In the last decade, these models have been extended to the domain of language by linguists and psycholinguists interested in linguistic categorization phenomena both in perception (e.g., Lacerda 1995; Goldinger 1996; Johnson 1997) and in production (Goldinger 2000; Pierrehumbert 2001).

In such models, each category is defined by a 'cloud' of remembered tokens, or exemplars, that have been associated with a category label. Exemplars are organized within the category by similarity across any salient dimension, producing internal structure in category-space; a given exemplar may therefore

contribute to many categories simultaneously. Depending on the model, new experiences are assigned to relevant categories by comparison to actual exemplars (Nosofsky 1986), or to generalizations emerging from the exemplars that make up a category (Hintzman 1986; Goldinger 1996; reviewed in Tenpenny 1995). Categorization of an experience is accompanied by insertion of a new, corresponding exemplar into category-space at the appropriate point, or by raising the activation of an indistinguishable exemplar that was previously stored (e.g., Kruschke 1992). Therefore, and crucially for the discussion that follows, each experience alters the entire category system slightly by changing the range and/or activation of component exemplars.

The finding that the phonetic details of recently perceived words resurface in a speaker's own productions of those words provides additional evidence for continuous updating of phonological categories through experience (Goldinger 2000). Most importantly in this context, this finding identifies a pathway for a production-perception feedback loop in adult speakers in which non-contrastive phonetic details of what is perceived can be subsequently reflected in the details of what is produced. (For some examples of modeling approaches to this production/perception loop, see Pierrehumbert 2001 and Oudeyer 2002.)

With the addition of the production/perception loop, exemplar models lend themselves particularly well to exploring the role of language change at the individual time-scale because they explicitly model a set of interacting factors that can result in Darwinian evolution:

- Variation: Categories are composed of many phonetically detailed variants; different speakers may have a distinct distribution of exemplars in a given category.
- Reproduction: Both production and perception are influenced by the exemplars within categories. Production details are influenced by the exemplars within the source category, and perception of an output token creates a new exemplar with corresponding details in memory.
- Selection: Any factor that influences the likelihood that a given exemplar will participate in production or that influences the way a given percept is likely to be categorized will influence the direction in which the category system updates over time. Exemplars that are more 'fit' by these criteria will leave a greater trace in the future behavior of the category than exemplars that are less fit.

An exemplar approach to modeling phonological categories thereby takes the opposite tack relative to many contemporary phonological theories which state that once acquired, phonological categories contain only highly abstract, non-redundant, distinctive information (discussed in Kenstowicz 1994). Within these theories production and perception make use of these phonological categories, but do not affect them. Conversely, within exemplar models each phonological category consists of a highly redundant, slowly updating set of many

phonetically detailed exemplars of that category, where the current state of the category influences its future behavior.

Exemplar models vary in their commitment to the exemplar metaphor of information storage. While it is possible, for example, to model categorization behavior as if all generalizations are abstracted as needed from an existing exemplar set, it is also possible to model the same behavior as proceeding through a sufficiently complex set of abstract generalizations that can be continually created and updated through experience (Tenpenny 1995). All of these variant models fulfill the basic criteria for natural selection laid out above, and so no stand will be taken here on whether exemplars represent the primary content of categories, or whether they instead serve as an intermediate buffer mediating between experience and continually updating, more abstract levels of generalization. For the purpose of computational simplicity however, in the examples that follow categories will be treated as constituted solely of exemplars, with no higher level of abstraction.

### **3. Exemplars and the propagation of variation**

In this article, I discuss three pathways that can influence the propagation of variation within a replicating population, and show how they may contribute to linguistic category evolution using an exemplar model as a heuristic tool. These pathways are (i) random fixation of variants through pruning of lines of descent, also known within the biological literature as *genetic drift*, (ii) blending inheritance, and (iii) natural selection itself. The first and third pathways are well attested mechanisms influencing change in biological populations, while the second, blending inheritance, is not.

These pathways have been chosen as illustrations because they may provide insight into several interesting issues in language change. First, given evidence that we perceive and imitate forms in a great deal of phonetic detail (Goldinger 2000), why does variation within categories fail to steadily increase with experience? Simple exemplar-based models of production-perception loops predict just such behavior, because all outputs alike are restored as new exemplars, regardless of error in production. In the absence of any mechanism to selectively remove outliers, this catholic storage behavior should eventually result in loss of distinction between categories as their exemplar distributions increasingly broaden and overlap (Pierrehumbert 2002). In Sections 3.1 and 3.2 below, we will see that pruning and blending inheritance are two plausible, non-exclusive mechanisms that can counteract creeping accumulation of variation even when forms can be perceived and imitated in great detail.

Second, how is phonological contrast maintained over the course of sound change? What mechanisms could support contrast phenomena such as vowel



chain shifts, or local anti-homophony effects? In Section 3.3, I will discuss evidence that lexical access processes may introduce selection against ambiguous pronunciations, and illustrate how this selective pressure may influence sound change in an exemplar-based model of an anti-homophony effect.

To better understand the workings of these mechanisms, we will start out the discussion with a fragmentary model of an exemplar-based production/perception loop. As the discussion proceeds, I will add complexity as appropriate. At the start, the model will consist of these five features:

1. Categories consist of detailed exemplars, each of which is associated with a base-activation level.
2. To model the greater influence of more recent memories, this base-activation level slowly decays over time (Nosofsky 1986; Pierrehumbert 2001).
3. Production from a category proceeds by creating a target based on a randomly chosen category exemplar, where the probability of choosing an exemplar is correlated with its base-activation level (Pierrehumbert 2001).
4. Resulting production tokens are stored as new exemplars in their category of origin, completing the production/perception loop.
5. These processes are subject to random noise, allowing a production token to deviate slightly from its exemplar basis, or a stored exemplar to deviate slightly from the source percept. This noise is the only source of new variation in the model.

### *3.1. Fixation of variants by pruning lines of descent*

Variation is the substrate for evolutionary change. Within an exemplar-based production/perception loop, if there is no noise (in the sense of stochastic variation) in what is produced or perceived, there is no way for a category to evolve beyond the range of the exemplars it starts with; change therefore depends on noise that introduces novel variants. However, recall that in a straightforward exemplar model, *every* perceived variant is stored, and potentially reproduced. If this is the case, it seems that noise-driven accumulation of variation should cause categories to inexorably broaden, eventually accumulating so much variation that they lose all contrast with other, nearby categories. In this section, I introduce a potential limit to such broadening through a biological example that can be straightforwardly illustrated in our simple model of language.

In an influential 1987 paper, Allan C. Wilson and colleagues presented evidence from a comparison of human mitochondrial DNA sequences that all humans share a female ancestor who lived roughly 200,000 years ago, probably in Africa (Cann et al. 1987). Mitochondria are subcellular organelles deriving from an ancient symbiosis between an early eukaryotic cell (the ancestor of modern animals, plants and fungi), and an ancestor of the *E. coli* bacterium.



Fitting with their origin as free-living bacteria, mitochondria retain a small number of their own genes, which replicate independently of the rest of the cell's nuclear DNA. In fertilization, the sperm only transfers nuclear DNA to the egg, with the result that the resulting fertilized egg contains only mitochondria from the mother. As a consequence, the line of descent for mitochondria (along with their DNA) is strictly matrilineal. If you are a man, your mitochondria will die with you, and if you are a woman, your mitochondrial line will have a chance to continue only if you have at least one daughter.

Mitochondrial DNA, just like nuclear DNA, slowly changes through random mutation. Any neutral mutation that appears in the mitochondrial DNA of a woman will appear in that of her daughters, and so each mutation can be thought of as initiating a unique line of descent. However, there are always women who fail to have daughters in any given generation, and in the rare event that *all* the women who carry a particular mitochondrial line fail to have daughters in some generation, that mitochondrial line of descent is lost, or 'pruned'. The larger the population that shares a particular line of descent, the less likely that a line will be pruned in a given generation, but the probability never reaches zero. By comparing the similarity of mitochondrial DNAs taken from people all over the world, and given well-supported assumptions about the rate of mutation, Wilson and colleagues were able to show that we all share a single mitochondrial line of descent that originated around 200,000 years ago. All other mitochondrial lines that originated at or before that time have been lost, and we are the more genetically homogeneous for it. The process of pruning, also known as 'genetic drift', can therefore be thought of as a mechanism that randomly eliminates variation within a set of self-reproducing elements. All of our mitochondria, in this example, are all related by vertical descent to those of a single woman, because the lines of descent leading toward the present from her contemporaries have all been lost (for further discussion, see Ayala 1995).

Just as lines of mitochondrial descent are pruned in biological reproduction, pruning reduces variation in our simple exemplar-based production-perception loop. Recall that for each production event, a single exemplar is chosen from a category at random as a model for a phonetic output. Each production event in turn results in a new exemplar stored in some category, which can itself give rise to new exemplars in the future. In this way, each current exemplar can trace its ancestry back in a line of descent. If we assume that the activation of exemplar memories decays slowly with time, and that probability of use of an exemplar as a model for production is proportional to its activation, then there will be an effective turnover of exemplars, as old exemplar memories die away and new ones take their place. Whenever an exemplar decays very far before it has a chance to be produced, its line of descent is effectively truncated. Given a finite number of exemplars in a category, and a long enough period of time,

eventually all but one line of descent will be pruned, such that at some point all exemplars are vertically related to the same ancestral exemplar. Note that pruning is a distinct process from natural selection: natural selection operates by affecting the reproductive rate of reproductive units on the basis of some trait, while pruning is simply a statistical effect of random loss of reproductive units before they reproduce. Depending on the selectional environment, natural selection can act to increase or decrease population variance of some trait, while pruning always results in a decrease in variance.

In this simple exemplar model, there is an additional factor encouraging loss of variation through pruning that does not have a well-accepted correlate in biological evolution<sup>2</sup>. Whenever an exemplar is used as the basis for a production token, it creates a new, highly activated exemplar somewhere, which is then itself likely to be used in turn, and so on. As a result, any particular lines of descent that gain an upper hand through random chance can rapidly become more highly populated at the expense of others. This positive feedback loop between use and future use is plausible, given experimental evidence that more recent exemplars have a greater influence on current production than older ones (Goldinger 2000).

Pruning of exemplar lines of descent is illustrated through a simulation based on our simple model (Figure 1). This simulation models the evolution of exemplar distributions within a single category where exemplars are allowed to vary along a single dimension in unit steps from 0–99. The simulation begins with the category primed with 10 exemplars at each unit value, i.e., with 10 exemplars each at each value from 0–99. Each beginning ‘line of descent’ group of 10 exemplars is associated in addition with a serial number from 0 to 99 that will be passed to all of their offspring, so that we can track the survival of lines of descent for each group. In each round of the simulation, four exemplars are randomly chosen in proportion to their activation to serve as the basis for a production token. The production token is subject to noise, with the result that it may deviate from its exemplar basis to a small degree. The token is then restored as a new exemplar, with the same serial number as its antecedent. Each new exemplar records in addition the cycle in which it was stored, and the base-activation of exemplars decreases exponentially as cycles proceed (after Pierrehumbert 2001). To help us identify loss of lines of descent, as well as for computational tractability, when the base-activation of an exemplar decays to the point that its probability of being chosen for a production event falls below .001 %, it is deleted from memory. For simplicity, the production/perception loop is modeled as operating over a single speaker/hearer;

---

2. Assortative mating, in which a trait is selected because it is favored by potential sexual partners, has been modeled in a way that does allow for this kind of positive feedback in biological evolution (Kondrashov and Kondrashov 1999; Dieckmann and Doebeli 1999).

results are not significantly different if exemplars are passed instead between multiple speaker/hearers (Wedel in prep).

Figure (1a) shows the distribution of exemplar values after 10,000 cycles, (1b) the remaining lines of descent after 10,000 cycles, and (1c) the rate of loss of lines of descent over the course of the simulation. In this typical simulation, at a production rate of only 4 exemplars per cycle, 95 % of the lines of descent have been lost at 10,000 cycles. As a result, despite continual addition of variation through noise, resulting in slowly broadening distributions about the original exemplar values (see Figure 1a), the total variation within the category has been starkly reduced. As this simulation shows, within an exemplar model that allows significant turnover in active exemplars over time, pruning of lines of descent may play a role in constraining the category broadening effects of noise.<sup>3</sup>

### 3.2. *Constraining variation through blending inheritance*

Blending inheritance, or character averaging, represents another variation-narrowing mechanism with quite interesting properties. Within the field of genetics prior to Mendel, it was assumed that an offspring's traits were an average of those of the parents. This 'blending inheritance' caused great difficulty for Darwin's theory of natural selection, because it would result in the steady elimination of variation within a population. This loss of variation would result because the traits of each successive generation, created as an average from the previous generation, must always fall between the extremes of the previous generation. If this steady averaging of traits really occurred, it was hard to understand how the observed level of variation could possibly persist in a population. Mendel's subsequent discovery of *particulate* inheritance, in which genes remain distinct in reproduction, showed how genetic variation could survive within offspring, even if covertly.<sup>4</sup>

3.2.1. *Within-category blending.* Inheritance within our current exemplar model in this article is particulate as well: each phonetic target is based on a single exemplar, and the characteristics of exemplars within a category do not influence the characteristics of any new exemplar as it is stored, and so, like

---

3. An additional example of a simulation of system change via pruning/genetic drift can be found at <http://ccl.northwestern.edu/netlogo/models/GenDriftPlocal>.

4. For example, consider the offspring of a parent with two genes for brown eyes, and a parent with two genes for blue eyes. These offspring will have brown eyes, because the brown-eye gene is dominant, but the blue-eye gene has not been lost. As a consequence, if a particular mating in later generations brings two blue-eye genes back together in some offspring, the blue-eye trait will resurface.




Figure 1. Simulation of pruning exemplar lines of descent

genes, our exemplars do not lose any of their characteristics in reproduction. However, just because genes are transmitted in close to a particulate manner, this does not mean exemplars or any other mental representations have to be. (See Sperber 1996 and Atran 2001 for arguments that mental representations are not particulate to the degree that genes are.) In a recent paper, Pierrehumbert (2002) addressed the problem of noise-driven category broadening by proposing that *multiple* exemplars within a region of a category contribute to the con-

struction of a phonetic target, rather than just one. If production targets are the result of averaging some number of exemplars within a region, then pre-noise production targets must always fall within the extremes defined by the exemplars within a category, providing a category-narrowing counterweight to the broadening influence of noise. (See also Hintzman 1986; Goldinger 1996 for additional examples of exemplar models that incorporate some form of blending inheritance to reduce variation.)

In the previous section, we saw that random pruning of lines of descent had the effect of diminishing the variation in categories, but did nothing to prevent exemplars in a category from splitting into widely separated subgroups, defined by distinct lines of descent (cf. Figure 1a). In contrast, Pierrehumbert's proposed mechanism of limiting variation introduces horizontal relationships between exemplars, in that exemplars from widely separated lineages within a category can cooperate to produce a production target. Provided no sub-groups of exemplars within a category are ever so distinct that the chance of cooperation between them falls to zero, some productions from a category will always fall between subgroups, acting to fill in gaps between them. Under this model, averaging between exemplars thereby produces a steady reversion to the mean over the entire category, such that in the absence of pressure to the contrary, all category exemplars will eventually fall into a unimodal distribution about the mean. To implement this here, I altered the simulation so that after choosing an exemplar at random from the category for production, an additional two exemplars were randomly chosen within a Gaussian window about the initial exemplar. These three exemplars were then averaged, in proportion to their respective activations, to produce a production token, which was then re-stored as a new exemplar in the category. Figure 2 shows the distribution of exemplars at cycle 1, 1000, and 5000 in a typical simulation, and here we see that even though the distribution of exemplars initially separates into multiple peaks, all peaks eventually coalesce into one.<sup>5</sup>

*3.2.2. Blending across categories.* Above, we saw that within-category blending inheritance over the production/perception cycle promotes development of a unimodal distribution of exemplars within that category. Are there any points within this cycle at which blending inheritance *across* phonological category boundaries might arise? Any stages of production and perception

---

5. Note that even though averaging between exemplars in production obscures vertical inheritance relationships, pruning is still operative: if any exemplar decays before being produced, the influence of its characteristics is lost to future generations. However, because characteristics are more diffused throughout in the population of exemplars in this model, and because they can arise through multiple pathways of averaging, pruning has relatively less ability to dramatically influence the distribution of exemplars in a category.




Figure 2. Within-category blending

that are likely to proceed without reference to phonological category boundaries would be good candidates for cross-category blending. In production for example, even though articulatory motor programs may be assembled under direction of the lexicon, the actual execution of those motor programs may proceed independently (Levelt 1989). Under a dynamical attractor model of motor control (Kelso et al. 1992; Kelso 1995; see also Saltzman and Munhall 1989), motor patterns form attractors which systematically bias motor output toward previously practiced outputs in relation to similarity (Zanone and Kelso 1992, 1997). If motor attractors deployed in production of distinct phonological categories can interact, nearby attractors will show a tendency to merge over time, in the absence of feedback to keep them apart, resulting in increasing overlap of the contents of neighboring phonological categories.

Within perception, another potential source of cross-category blending lies in the perceptual magnet effect (Kuhl 1991). The perceptual magnet effect refers to the finding that perceptual space is warped toward category centers relative to objective stimuli dimensions, such that percepts near category centers are perceived as closer together than they actually are, while percepts near category boundaries are perceived as farther apart. In exemplar terms, new exemplars will not correspond straightforwardly to the stimuli that gave rise to them, but will be biased systematically toward the centers of existing exemplar distributions.

Within the model of Guenther and Gjaja (1996), the perceptual magnet effect arises early in perception through non-linearities in the distribution of firing preferences within cortical maps tuned to auditory stimuli. Within this model, the cortical map develops in response to the statistical distribution of sounds perceived by an infant and does not require higher level phonological categorization (cf. the earlier models of Kuhl 1991, 1995; Lacerda 1995, in which categorization precedes warping). Given evidence that established cortical maps

retain a significant degree of plasticity in adulthood, that is, can continuously shift in response to experience (Weinberger 2003; Guenther et al. 2004; Eisner and McQueen 2005; Xerri et al. 2005; Noreña and Eggermont 2005), changes in the distribution of inputs to the perceptual system should slowly provoke shifts in the pattern of perceptual warping (Case et al. 1995). Any alteration in the pattern of warping will necessarily influence the characteristics of new exemplars feeding into phonological categories, and within our model, therefore influence the characteristics of production tokens proceeding from those categories.

If biasing of percepts occurs prior to categorization, as in the model of Guenther and Gjaja, then the perceptual magnet effect creates a dimension-wide blending inheritance in perception: given an exemplar distribution along some dimension, percepts will be, on average, warped away from the extremes toward the center, because there will be, on average, more exemplars toward the center of a distribution than toward the edges. The result is a steady reversion toward the mean of the entire distribution in perception.<sup>6</sup>

However, we know that although phonological categories occasionally merge (as in the [ɔ ~ a] merger in the western USA), phonological distinctions do not all inexorably collapse, as cross-category blending would dictate. In fact, categories frequently evolve as if they were actively avoiding merger, as in cases of push- and pull-chain vowel shifts (Labov 1994: Chapters 5, 6). Within our model so far however, there is no way to model this: if there are any points in the production/perception loop in which cross-category blending inheritance is significant, category merger should be the rule.

### 3.3. *Natural selection and evolution of phonological categories*

What reasons do we have for thinking that there may in fact be some factor(s) in language change that would mitigate such a tendency for categories to collapse? Given that information transmission requires contrast within the signal, a long literature in linguistics has proposed that mechanisms must exist to promote the maintainance of contrastive units that constitute the speech signal (e.g., Gillieron 1918; Martinet 1955; and many more).

---

6. This is true even if the exemplar distribution along a dimension divided into multiple peaks. Within the model of Guenther and Gjaja (1996), a stimulus is perceived as the population vector of *all* neurons responding to the stimulus, and as a consequence, a stimulus falling between two response peaks in the cortical map will experience less average warping toward either peak than a stimulus falling a similar distance away from only one response peak. Therefore, even though stimuli falling near the boundary of two perceptual categories produce percepts that are warped away from the boundary as per the perceptual magnet effect, the *average* of all percepts in each category will shift over time toward the boundary.



For example, many linguists have considered the hypothesis that the probability of a global merger between two phonemes may be inversely correlated with the functional load of their opposition, that is, the degree to which lexical items depend on their opposition for contrast (Martinet 1955). As functional load is usually proposed to be only one of many factors influencing the probability of a merger (e.g., Labov 1994: 328), this hypothesis cannot be accepted or rejected on the basis of analysis of one or two cases, but requires the statistical analysis of a large sample of mergers, where the relative functional loads of the pre-merger opposition and other non-merged oppositions can be accurately measured and compared (King 1967; Surendran and Niyogi 2003). This is a non-trivial enterprise, and a comparison of sufficient scale to provide a strong test of this hypothesis has yet to be carried out (Surendran and Niyogi 2003).

Two conceptually related hypotheses that may be easier to approach are (i) that an otherwise regular sound change can be locally inhibited where it would eliminate a crucial contrast, and (ii) that novel mechanisms of contrast can develop locally in response to low functional contrast between lexical items. As an example of the first type, Juliette Blevins (2005) argues that in a majority of cases, synchronic antigemination effects are not phonetically natural, but occur where otherwise regular gemination would result in the obliteration of a paradigmatic contrast. In Iraqi Arabic, for example (Erwin 1963: 56–58, as cited in Blevins 2005), a syncope rule deletes short vowels from stem-final syllables when a vowel-initial suffix is added. Iraqi Arabic also has a degemination rule that converts non-prevocalic geminates into singletons. Vowel syncope can feed degemination, as shown in the example in (1a) below. However, syncope does not occur in inflected verbs of certain classes, precisely where syncope plus degemination would produce an output homophonous with another form in the paradigm, as in (1b).

- |     |    |                                 |        |  |
|-----|----|---------------------------------|--------|--|
| (1) | a. | <i>ybaddil</i><br>'he changes'  | versus | <i>ybadluun</i> < /y <b>b</b> addil-uun/<br>'they change'  |
|     | b. | <i>tammam</i><br>'he completed' | versus | <i>tammimaw</i> < / tammam-aw/<br>'they completed'<br>(* <i>tammaw</i> , cf. <i>tammaw</i><br>'they are complete') |

While it is clearly not the case that incipient loss of contrast always prevents a sound change from proceeding locally, at least in the case of anti-gemination, the converse correlation appears to be strong: Blevins argues from a broad survey of data that anti-gemination effects are rare except when gemination would result in some form of intra-paradigmatic contrast loss, suggesting a causal relationship. Additional discussion of anti-homophony effects in paradigms can be found in Yip (1998), Crosswhite (1999) and Gessner and Hansson (in press).

The second hypothesis springs from the observation that when a contrastive feature is lost in some context through sound change, a previously non-contrastive, phonetic variation conditioned by the erstwhile feature can become phonologized, in what looks tantalizingly like compensation (Kiparsky 1932; Twadde 1938, discussed in Janda 2003). Examples include the emergence of contrastive vowel nasalization in concert with the loss of coda nasals in French, or the development in German of the front, round vowels /y/ and /ø/ from /u/ and /o/, in step with the loss of a conditioning /i/ or /j/ in a following syllable. Two issues that arise from this observation will be of interest to us here. First, it has long been noted that the transition of a phonetic variant from the status of non-contrastive to contrastive must be gradual, beginning before the full loss of the conditioning contrast. Otherwise, at the point when the conditioning feature is lost, the conditioned phonetic variant would cease to exist as well (discussed in Janda 2003: 409). Second is the issue of mechanism. While the emergence of a new contrast at the same time that an old one is lost would seem to serve the larger functional purpose of maintaining sufficient contrast in the system of communication, this observation provides no causal hypothesis. Reacting against teleological arguments for sound change in general, Ohala in fact rejects the notion that contrast plays any role in this phenomenon, arguing instead that ‘hypocorrection’, that is, underapplication of compensatory perception rules, is solely responsible for phonologization of allophonic variants in this context (Ohala 1989).

Here, I would like to argue that evolutionary, selectional processes can also provide us with a gradual, non-teleological mechanism for this phenomenon which retains an indirect causal connection to contrast maintenance. Up to this point, the production/perception loop in our model has been unbiased: the odds that any exemplar will contribute to producing a token over its lifetime are even, and every token produced results in the storage of a new exemplar in some category. If, however, any factor intervenes to alter the likelihood that a given property will be transmitted successfully within the production/perception loop, the frequency of this property will change accordingly over time. There are many points at which biases could enter. In production for example, contextual usage biases toward lenition (reviewed in Bybee 2001) or hyperarticulation (e.g., Lindblom 1986; Johnson et al. 1993) could shift the details of particular lexical categories away from the mean of their component phoneme categories over time. At the point of categorization on the other hand, any factor that influences the likelihood that a percept will be successfully categorized will in turn influence the distribution of incoming exemplars. In the present context, I will concentrate on the latter possibility to show how it can mitigate the influence of cross-category blending inheritance to preserve category contrast.

A great deal of work on factors influencing lexical access has shown that lexical identification is less accurate for words that have many similar lexical

neighbors (Luce 1986; Goldinger et al. 1989; Cluff and Luce 1990; reviewed in Luce and Pisoni 1998). Under the Neighborhood Activation Model (NAM, Luce and Pisoni 1998), a stimulus (i.e., a spoken word) activates all lexical categories in the lexicon in relation to similarity, which then compete for recognition. Stimuli that strongly activate only one lexical category are processed efficiently because there is little competition, but stimuli that simultaneously strongly activate competing lexical categories will be processed more slowly, with greater opportunity for error. NAM further predicts that if activation levels of competing lexical categories are sufficiently well-matched, processing may be delayed to the point that the percept will fail to be assigned to any lexical category at all (Luce and Pisoni 1998). This suggests that when a pronunciation strongly activates more than one lexical category, it has a greater chance of failing to be identified altogether. This hypothesis is supported by the finding that within a lexical decision task, real words with high frequency neighbors are more often mistakenly reported as nonwords than real words without high frequency neighbors (Luce and Pisoni 1998).

Under our exemplar model of the production loop, such unidentified pronunciations will fail to contribute a new exemplar to the system, and therefore fail to contribute to the evolution of lexical categories. As a consequence, the exemplar characteristics that give rise to more ambiguous pronunciations will have lower overall reproductive fitness in the sense that they contribute less to future pronunciations, and so should tend to decrease in frequency over time. In this model, shifts in the burden of contrast are encouraged when an incipient loss of contrast in one part of the system makes lexical access less efficient for the near homophones that depend on that particular contrast. The resulting selection pressure against more ambiguous pronunciations of these near homophones within a production/perception loop allows any conditioned contrast variation elsewhere in the word to be amplified, converting what was originally low-level, gradient variation into more categorical difference.

To provide a more concrete example, imagine that we have two close lexical neighbors that would traditionally be understood as abstractly differing by one distinctive feature, say [kæt] and [kæd]. Within an exemplar model of the lexicon, where lexical entries are represented in terms of detailed phonetic distributions, the average distributions of phonetic properties representing each of our two lexical items will differ not only in the mean degree of voicing of the final stop, but also in any allophonic variation associated with differences in stop voicing. In this case, for example, we can expect that tokens of [kæd] will often have a relatively longer preceding vowel than we find in tokens of [kæt], as a phonetic effect of the following voiced stop. Listeners have been shown to be very sensitive to non-contrastive allophonic detail in identifying percepts (e.g., Andruski et al. 1994; Gow and Gordon 1995; McMurray et al. 2002), and so we expect that this average vowel length difference would contribute to distin-

guishing between production tokens of these two words. In our example then, if a token of [kæd] happens not to be well distinguished from [kæt] by its degree of final-voicing, but still has a relatively long vowel, it will be more likely to be understood and imitated than a similar token of [kæd] with a vowel of intermediate length. As a consequence, any consistent bias toward neutralizing the stop voicing distinction between [kæt] and [kæd] in a production/perception loop will provide an opportunity for a gradual shift in the burden of contrast from final stop voicing to vowel length. If we grant cognitive or motor mechanisms of generalization that can generate global behaviors from local patterns (e.g., Bybee 2002), word-based shifts such as this could form the seed pattern for a more general pattern change. Continuing the example above, if the shift in contrast-load from final obstruent voicing to the length of the preceding vowel proceeds sufficiently in close neighbor pairs such as [kæt] and [kæd], pattern generalizing properties of the production and/or perception mechanism could extend it to other words lacking a correspondingly close neighbor. In this way, selection-driven pressure to maintain local lexical contrasts represents a plausible, non-teleological pathway contributing to the maintenance of systemic contrasts.

To illustrate this selection-driven model of contrast development, we modify our simulation to include lexical as well as segment categories, and to include competition between lexical categories for percepts (Wedel 2004, in prep.). We will start out with two lexical categories, A and B, where lexical exemplars are composed of two ordered segment exemplars that can vary within the range 0–20. A possible lexical exemplar in this simulation would be [[10], [5]], for example, where the first component segment exemplar has the value ten, and the second the value five. Square brackets enclose exemplars (see Pierrehumbert 2003 for discussion of nesting of exemplar-based categories). Construction of a phonetic target for a category proceeds by randomly choosing a lexical exemplar from that category in relation to its base activation, and then averaging each of the component segment exemplars with two other nearby segment exemplars as before, producing within-category blending inheritance. Mimicking the perceptual magnet effect as modeled by Guenther and Gjaja (1996), cross-category blending proceeds by warping of percepts towards distributional maxima in the entire distribution of previously perceived segments. This is done by calculating a ‘population vector’ (Georgopolis et al. 1983, 1988) over the current activations of all segment exemplars relative to the current percept, and averaging the percept with that vector.

Identification of a warped percept with a category follows Nosofsky’s exemplar-based generalized context model for categorization (GCM; Nosofsky 1986, 1988). GCM is conceptually parallel to NAM, but in addition models the influence of intra-category variation on categorization behavior by including exemplar-based storage. Upon presentation with a warped percept, lexical

exemplars in all categories become activated above their base level in proportion to their similarity with the percept, where similarity is calculated using an exponential decay function of the Euclidean distance between percept and exemplar (Nosofsky 1986; see also Kruschke 1992). The percept is probabilistically identified with a category in proportion to the summed activation of its component exemplars following Luce's choice rule (Luce 1959), and stored there as a new exemplar.

To include within the simulation a feature modeling the finding that ambiguous tokens are less fit in categorization, the simulation is modified to allow occasional, stochastic failure of categorization, where the probability of failure is greater to the degree that exemplar activation is distributed among multiple lexical categories. In the example simulation that follows, this probability (and therefore the possible difference in fitness between tokens) was kept very low. At the lowest possible contrast, that is, in the limiting case in which a token activated all categories equally, the rate of failure was set at 10%; with greater degrees of difference in categorical activation, the rate of failure was correspondingly lower.

Previous simulations including these features have demonstrated, as expected, that the small selection pressure against ambiguous production tokens results in a tendency for lexical categories to retain contrast, despite leveling pressure from cross-category blending inheritance (Wedel 2004, in prep.). Here, I show a simulation using our current model which illustrates a potential role for covert variation in selection-driven evolution of contrast shifts under neutralizing change.

The simulation begins with two lexical categories, A and B, each populated with 100 identical exemplars of the following form: A: [[10], [10]], and B: [[10], [0]]. To mimic the variation-inducing effect of contextual conditioning, low values of a final segment were made to increase the probability of a higher value in the previous segment. Specifically, final segment values below 5 created a slight bias in the production of the preceding segment toward output values of at least 10. As a consequence, given the second-segment value of 0 in the exemplars in category B, production tokens from category B are more likely to surface with initial-segment values above 10 than tokens from category A, which have higher second-segment values.

Figure 3 shows the average exemplar value for each of the two segments for each of the two categories over 2000 cycles, labeled A-1, A-2 and B-1, and B2 respectively. Noise in production ensures that each category rapidly relaxes from the initial uniform state to become composed of a distribution of exemplars, but at the same time, intra-and inter-category blending in production ensure that distributions do not expand to occupy the entire provided space. At the beginning of the simulation, the lexical exemplars in the two lexical categories are distinct only in their final segment values, of 10 versus 0. Given the




Figure 3. Change in average segmental exemplar values in categories A and B over 2000 cycles

parameter settings of the simulation used here, this degree of contrast is sufficient to reliably distinguish output tokens proceeding from each category, with the result that the system is stable over many cycles (Figure 3, cycles 0–500). Note as well that the average value of the first segment of category B is slightly higher than that of category A, due to the conditioning bias of the following low-value segment in B exemplars.

At cycle 500, to simulate a neutralizing sound change in progress, I introduce an additional slight bias into the production of all output tokens toward reduction of the value of all segments in final position. As a consequence, the average final segment value in production tokens proceeding from category A (labeled A-2 in Figure 3) starts to slowly trend downward from the original value near 10 (Figure 3, cycles 500–2000). As this segment shifts its average value downward, output tokens from categories A and B become increasingly confusable. The typical standard deviation of the sound exemplars in each lexical category is on the order of 4–5 units, so by cycle 1200, more and more of the production tokens originating from categories A and B are similar to exemplars in both categories and are therefore stored at a slightly lower rate. After cycle 1200, however, the average value of the first segment in B (labeled B-1 in Figure 3) begins to rise, becoming increasingly distinct from the first segment of A. Within a few hundred cycles, this difference becomes entrenched, and then the average values of the final segments jointly coalesce near zero.

Why do these changes occur in this structured way? First, recall that noise produces a constant supply of low-frequency segment variants. When fitter variants happen to arise in the context of some selection pressure, they can gain a foothold and ‘found’ a new line of descent that eventually takes over the category. A representative portion of this process can be seen in Figure (4), which shows the distribution of segment exemplars in the second position of




Figure 4. Change in the distribution of first-segment exemplar values in category B from cycle 1220 to 1340

B from cycles 1220–1340. In this particular simulation, by around cycle 1200 a small but significant number of segment exemplars around the values 14–17 had arisen in the first position of B, although the average was still near 12. The more of these exemplars with higher first-segment value there were in B, the more likely similar production tokens were to be stored there, producing in turn more production tokens from B with high values in the first position, and so on. By cycle 1260, as this process began to snowball, the distribution of segment exemplars in the first position of B became strongly bimodal, with an older set of exemplars with first-position segment values near 12, and a largely newer set of exemplars with first-position values nearer to 17. By cycle 1340, the shift to the higher value was nearly complete, providing the two categories a means for contrast beyond the difference between the final segments.

In this model, natural selection at the level of categorization is the underlying process driving conversion of a gradient, ‘phonetic’ distinction into a more nearly categorical contrast as an associated contrast is neutralized. In this particular simple simulation, the original small difference in average values in the first segments of A and B exemplars derived from a small production bias imposed by a following, low-value segment. This difference remained small because of the variance-reducing effects of intra-category blending. However, as production tokens from the two categories became increasingly similar as the two final segments approached one another, the selectional bias against ambiguous tokens caused the initially small difference in first segment values to expand, until the two segments became sufficiently different to be contrastive within the confines of the simulation.



Of course, this highly simplified simulation does not include a great number of factors that are relevant to real-world sound change, but I hope it has succeeded in illustrating how an increase in the rate of lexical access failure for ambiguous tokens can result in evolution of greater contrast between lexical categories. A property of this type of model is that the greatest selection pressure falls on near-homophones as a neutralization is in progress, with the result that the compensatory change is predicted to advance fastest here. As the change becomes consolidated in near-homophones, other words that share sound sequence exemplars with the words leading the change may follow suit (Bybee 2001, 2002). This latter prediction highlights a further feature of models in which exemplar categories exist at both the word, and sound/sound sequence levels, creating a densely interconnected web within the lexicon: change at the word level can feed back to the entire system, allowing a sufficiently widespread word-centered change to initiate global, regular change (Pierrehumbert 2002; Bybee 2002; Wedel 2004).

To the extent that change in speech patterns over the lifetime of an individual can significantly influence language change (see, e.g., Labov 1994: Ch. 2), language change must be thought of as the result of two nested cycles: a smaller scale cycle driven by language use by speakers, and the longer scale cycle of transmission between generations. Although I have modeled change in the simulations presented here as occurring within a single individual, many of these changes, such as the global entrenchment of novel contrasts, may in fact require the discontinuity provided by transmission to children (discussed in, e.g., Lightfoot 2003; Blevins 2004). An additional step toward realism could be taken by modeling more modest shifts in category contents within individuals over time by mechanisms such as pruning, blending and selection, superimposed on multiple cycles of transmission to naïve speakers (e.g., Hare and Elman 1995; Kirby and Hurford 2002).

#### **4. Conclusions**

Exemplar models arose in response to widely supported findings that people store a great deal specific detail in memory, as well as arbitrary associations between details. This rich memory for associations suggests that people can store details in the form of their source episodes, without the degree of information loss that would accompany sifting out pre-established category features, or converting a percept into a prototype. The resulting particulate view of category contents allows the direct application of evolution theory in modeling category change, because all variation in categorized percepts is potentially maintained, even when that variation is not relevant to the processing of the category in its present state. As in the evolution of genetic systems, within an

exemplar-based production-perception loop low-level variation in the contents of a category serves as a reserve of seed patterns for shifts in the system over time.

Patterns of evolution in biological systems serve as suggestive examples for hypothesizing about pathways of language change. Nonetheless, some evolutionary patterns of change in language are likely to be different from common patterns in biology in important ways because the infrastructures storing and expressing patterns are distinct in the two kinds of systems. Biological patterns are encoded in genes and expressed in bodies, while cognitive patterns are encoded in neuronal structures and expressed in thought and behavior. Here, we explored the ability of an exemplar model to illuminate two potential pathways for linguistic category change that have conceptual parallels in biological evolution: change through pruning of lines of descent, and change through natural selection. In contrast, the pathway of blending inheritance does not play a large role in biological evolution because genes are transmitted to a large degree discretely. However, to the extent that cognitive representations are non-discrete, it is likely that blending inheritance may play a significant role in the evolution of linguistic categories. This article for one adds to a growing body of theoretical research that identifies blending inheritance as a potentially significant mechanism to maintain category cohesion against the steady introduction of variation through experience (see, e.g., Hintzman 1986; Goldinger 1996; Pierrehumbert 2002; Wedel 2004). As additional research is done on the character of cognitive representations and their interaction, our understanding of these cognition-specific mechanisms of evolutionary change will continue to develop.

*University of Arizona*

## References

- Andruski, Jean E., Sheila E. Blumstein and Martha W. Burton (1994). The effect of subphonetic differences on lexical access. *Cognition* 52: 163–187.
- Atran, Scott (2001). The trouble with memes: inference versus imitation in cultural creation. *Human Nature* 12: 351–381.
- Ayala, Francisco J. (1995). The myth of Eve, molecular biology and human origins. *Science* 270: 1930–1937.
- Baudouin de Courtenay, Jan ([1895] 1972). An attempt at a theory of phonetic alternations. In *A Baudouin de Courtenay Anthology. The Beginnings of Structural Linguistics*, Edward Stankiewicz (ed.), 144–212. Bloomington: Indiana University Press.
- Blevins, Juliette (2004). *Evolutionary Phonology. The Emergence of Sound Patterns*. Cambridge, UK: Cambridge University Press.
- (2005). Understanding antigemination. In *Linguistic Diversity and Language Theories*, Zygmunt Frajzyngier (ed.), 203–234. Amsterdam/Philadelphia: John Benjamins.
- Bybee, Joan (2001). *Phonology and Language Use*. Cambridge: Cambridge University Press.

- (2002). Word frequency and context of use in the lexical diffusion of phonetically conditioned sound change. *Language Variation and Change* 14: 261–290.
- Cann, Rebecca L., Mark Stoneking and Allan C. Wilson (1987). Mitochondrial DNA and human evolution. *Nature* 325: 31–36.
- Case, Pamela, Betty Tuller, Mingzhou Ding and J. A. Scott Kelso (1995). Evaluation of a dynamical model of speech perception. *Perception and Psychophysics* 57: 977–988.
- Cluff, Michael S., and Paul A. Luce (1990). Similarity neighborhoods of spoken two-syllable words. Retroactive effects on multiple activation. *Journal of Experimental Psychology, Human Perception & Performance* 16: 551–563.
- Cooper, David L. (1999). *Linguistic Attractors: The Cognitive Dynamics of Language Acquisition and Change*. Amsterdam/Philadelphia: John Benjamins.
- Crosswhite, Katherine (1999). Inter-paradigmatic homophony avoidance in two dialects of Slavic. *UCLA Working Papers in Linguistics 1, Papers in Phonology* 2: 48–67.
- Cziko, Gary (1995). *Without Miracles. Universal Selection Theory and the Second Darwinian Revolution*. Cambridge, MA: MIT Press.
- Dawkins, Richard (1982). *The Extended Phenotype*. Oxford: Oxford University Press.
- (1983). Universal Darwinism. In *Evolution, From Molecules to Men*, D. S. Bendall (ed.), 403–425. Cambridge: Cambridge University Press.
- Dennet, Daniel C. (1995). *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. New York: Simon and Schuster.
- Dieckmann, Ulf and Michael Doebeli (1999). On the origin of species by sympatric speciation. *Nature* 400: 354–356.
- Eisner, Frank and James M. McQueen (2005). The specificity of perceptual learning in speech processing. *Perception and Psychophysics* 67: 224–238.
- Erwin, Wallace M. (1963). *A Short Reference Grammar of Iraqi Arabic*. Washington, DC: Georgetown University Press.
- Georgopoulos, Apostolos P., Roberto Caminiti, John F. Kalaska and Joseph T. Massey (1983). Spatial coding of movement: A hypothesis concerning coding of movement direction by motor cortical populations. *Experimental Brain Research. Supplement* 7: 327–336.
- Georgopoulos, Apostolos P., Richard Kettner and Andrew Schwartz (1988). Primate motor cortex and free arm movement to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population. *Journal of Neurosciences* 8: 2928–2937.
- Gessner, Suzanne and Gunnar O. Hansson (in press). Anti-homophony effects in Dakelh Carrier valence morphology. *Proceedings of BLS 30*. Berkeley: Berkeley Linguistics Society.
- Gillieron, Jules (1918). *Genealogie des mots qui designent l'abeille*. Paris: Champion.
- Goldinger, Stephen D. (1996). Words and voices, Episodic traces in spoken word identification and recognition memory. *Journal of Experimental Psychology, Learning Memory and Cognition* 22: 1166–1182.
- (2000). The role of perceptual episodes in lexical processing. In *Proceedings of SWAP, Spoken Word Access Processes*, Anne Cutler, James M. McQueen and Rian Zondervan (eds.), 155–159. Nijmegen: Max-Planck-Institute for Psycholinguistics.
- Goldinger, Stephen D., Paul A. Luce and David B. Pisoni (1989). Priming lexical neighbors of spoken words. Effects of competition and inhibition. *Journal of Memory and Language* 28: 501–518.
- Gow, David and Peter Gordon (1995). Lexical and prelexical influences on word segmentation. Evidence from priming. *Journal of Experimental Psychology, Human Perception and Performance* 21: 344–359.
- Guenther, Frank H. and Marin N. Gjaja (1996). The perceptual magnet effect as an emergent property of neural map formation. *Journal of the Acoustical Society of America* 100: 1111–1121.

- Guenther, Frank H., Alfonso Nieto-Castanon, Satrajit S. Ghosh and Jason A. Tourville (2004). Representation of sound categories in auditory cortical maps. *Journal of Speech, Language, and Hearing Research* 47: 46–57.
- Hare, Mary and Jeffrey L. Elman (1995). Learning and morphological change. *Cognition* 56: 61–98.
- Harrington, Jonathan, Sallyanne Palethorpe and Catherine I. Watson (2000). Does the Queen speak the Queen's English? *Nature* 408: 927–928.
- Henrich, Joseph and Richard Boyd (2002). On modeling culture and cognition: Why cultural evolution does not require replication of representations. *Culture and Cognition* 2: 87–112.
- Henrich, Joseph, Richard Boyd and Peter J. Richerson (in press). Five misunderstandings about cultural evolution. In *Epidemiology of Ideas*, Daniel Sperber (ed.). Chicago/La Salle: Open Court Publishing.
- Hintzman, Douglas L. (1986). "Schema abstraction" in a multiple-trace memory model. *Psychological Review* 93: 411–428.
- Hurford, James (1999). The evolution of language and languages. In *The Evolution of Culture*, Robin Dunbar, Christopher Knight and Carol Power (eds.), 173–193. Edinburgh: Edinburgh University Press.
- Jacoby, Larry L. and Lee Brooks (1984). Nonanalytic cognition, memory, perception and concept learning. In *The Psychology of Learning and Motivation* (Advances in Research and Theory, 18), Gordon H. Bower (ed.), 1–47. San Diego, CA: Academic Press.
- Janda, Richard D. (2003). 'Phonologization' as the start of dephoneticization – or, on sound change and its aftermath, of extension, generalization, lexicalization and morphologization. In *The Handbook of Historical Linguistics*, Richard D. Janda and Brian D. Joseph (eds.), 401–422. Oxford: Blackwell.
- Johnson, Keith (1997). Speech perception without speaker normalization. In *Talker Variability in Speech Processing*, Keith Johnson and John W. Mullennix (eds.), 145–166. San Diego: Academic Press.
- Johnson, Keith, Edward Flemming and Richard Wright (1993). The hyperspace effect: Phonetic targets are hyperarticulated. *Language* 69: 505–528.
- Kelso, J. A. Scott (1995). *Dynamic Patterns. The Self-Organization of Brain and Behavior*. Cambridge, MA: MIT Press.
- Kelso, J. A. Scott, Mingzhou Ding and Gregor Schöner (1992). Dynamic pattern formation: A primer. In *Principles of Organization in Organisms*, Arthur B. Baskin and Jay Mittenthal (eds.), 397–439. Santa Fe Institute, Santa Fe, NM: Addison-Wesley Publishing Co.
- Kenstowicz, Michael (1994). *Phonology in Generative Grammar*. Cambridge, MA and Oxford: Blackwell.
- King, Robert (1967). Functional load and sound change. *Language* 43: 831–852.
- Kiparsky, Valentin (1932). Johdatusta fonologiaan. *Virittäjä* 36: 230–250.
- Kirby, Simon and James Hurford (2002). The emergence of linguistic structure. An overview of the iterated learning model. In *Simulating the Evolution of Language*, Angelo Cangelosi and Dominic Parisi (eds.), 121–148. London: Springer.
- Kondrashov, Alexey S. and Fyodor A. Kondrashov (1999). Interactions among quantitative traits in the course of sympatric speciation. *Nature* 400: 351–354.
- Kruschke, John. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review* 99: 22–44.
- Kuhl, Patricia K. (1991). Human adults and human infants show a 'perceptual magnet effect' for the prototypes of speech categories, monkeys do not. *Perception and Psychophysics* 50: 93–107.
- (1995). Mechanisms of developmental change in speech and language. In *Proceedings of the XIIIth International Congress of Phonetic Sciences, Vol. 2.*, Kjell Elenius and Peter Branderud (eds.), 132–139. Stockholm: Stockholm University.
- Labov, William (1994). *Principles of Language Change. Internal Factors*. Oxford: Blackwell.

- Lacerda, Franciso (1995). The perceptual-magnet effect. An emergent consequence of exemplar-based phonetic memory. In *Proceedings of the XIIIth International Congress of Phonetic Sciences, Vol. 2.*, Kjell Elenius and Peter Branderud (eds.), 140–147. Stockholm: Stockholm University.
- Levelt, Willem J. M. (1989). *Speaking*. Cambridge MA: MIT Press.
- Lewontin, Richard C. (2005). The wars over evolution. *New York Review of Books* 52 (16): 51–54.
- Lightfoot, David (2003). Grammatical approaches to syntactic change. In *The Handbook of Historical Linguistics*, Richard D. Janda and Brian D. Joseph (eds.), 496–508. Oxford: Blackwell.
- Lindblom, Björn (1986). Economy of speech gestures. In *The Production of Speech*, Peter MacNeilage (ed.), 217–245. New York: Springer-Verlag.
- (1992). Phonological units as adaptive emergents of lexical development. In *Phonological Development. Models, Research, Implications*, Charles A. Ferguson, Lise Menn and Carol Stoel-Gammon (eds.), 131–163. Timonium: York Press.
- Luce, Paul A. (1986). Neighborhoods of words in the mental lexicon. *Research on Speech Perception, Tech. Rep. No. 6*, Bloomington, Indiana.
- Luce, Paul A. and David B. Pisoni (1998). Recognizing spoken words, The neighborhood activation model. *Ear and Hearing* 19: 1–36.
- Luce, R. Duncan (1959). *Individual Choice Behavior*. New York: Wiley.
- Martinet, André (1955). *Economie des changements phonétiques*. Bern: Francke.
- McMurray, Bob, Michael K. Tanenhaus and Richard N. Aslin (2002). Gradient effects of within-category phonetic variation on lexical access. *Cognition* 86: B33–B42.
- Noreña, Arnaud J. and Jos J. Eggermont (2005). Enriched acoustic environment after noise trauma reduces hearing loss and prevents cortical map reorganization. *The Journal of Neuroscience* 25: 699–705.
- Norris, Dennis, James McQueen and Anne Cutler (2003). Perceptual learning in speech. *Cognitive Psychology* 47: 204–238.
- Nosofsky, Robert (1986). Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology, General* 115: 39–57.
- (1988). Similarity, frequency, and category representations. *Journal of Experimental Psychology, Learning, Memory and Cognition* 14: 54–65.
- Oghala, John J. (1989). Sound change is drawn from a pool of synchronic variation. In *Language Change. Contributions to the Study of its Causes* (Trends in Linguistics, Studies and Monographs No. 43), Leiv E. Breivik and Ernst H. Jahr (eds.), 173–198. Berlin: Mouton de Gruyter.
- (1992). What's cognitive, what's not, in sound change. In *Diachrony within Synchrony, Language History and Cognition* (Duisburger Arbeiten zur Sprache und Kulturwissenschaft 14), Guenter Kellermann and Michael D. Morrissey (eds.), 309–355. Frankfurt am Main: Peter Lang Verlag.
- Oudeyer, Pierre-Yves (2002). Phonemic coding might be a result of sensory-motor coupling dynamics. In *Proceedings of the 7th International Conference on the Simulation of Adaptive Behavior*, Bridget Hallam, Dario Floreano, John Hallam, Gillian Hayes and Jean-Arcady Meyer (eds.), 406–416. Cambridge: MIT Press.
- Paul, Hermann (1886). *Prinzipien der Sprachgeschichte*. Halle: Niemeyer.
- Pierrehumbert, Janet (2001). Exemplar dynamics, word frequency, lenition, and contrast. In *Frequency Effects and the Emergence of Linguistic Structure*, Joan Bybee and Paul Hopper (eds.), 137–157. Amsterdam: John Benjamins.
- (2002). Word-specific phonetics. In *Laboratory Phonology VII*, Carlos Gussenhoven, Tanya Rietvelt and Natasha Warner (eds.), 101–139. Berlin/New York: Mouton de Gruyter.
- (2003). Phonetic diversity, statistical learning, and acquisition of phonology. *Language and Speech* 46: 115–154.
- Saltzman, Elliot and Kevin G. Munhall (1989). A dynamical approach to gestural patterning in speech production. *Ecological Psychology* 1: 333–382.

- Sancier, Michelle L. and Carol A. Fowler (1997). Gestural drift in a bilingual speaker of Brazilian Portuguese and English, *Journal of Phonetics* 25: 421–436.
- Sperber, Daniel (1996). *Explaining Culture. A Naturalistic Approach*. Oxford, UK: Cambridge, MA: Blackwell.
- Surendran, Dinoj and Partha Niyogi (2003). Measuring the usefulness (functional load) of phonological contrasts. *Technical Report TR-2003-12*, Department of Computer Science, University of Chicago.
- Tenpenny, Patricia L. (1995). Abstractionist versus episodic theories of repetition priming and word identification. *Psychonomic Bulletin and Review* 2: 339–363.
- Twaddell, William (1938). A note on Old High German umlaut. *Monatshefte für deutschen Unterricht* 30: 177–181.
- Wedel, Andrew B. (2004). Self-organization and categorical behavior in phonology. PhD dissertation, University of California, Santa Cruz.
- (in preparation). Disruptive selection and contrast maintenance in language change.
- Weinberger, Norman M. (2003). The nucleus basalis and memory codes: Auditory cortical plasticity and the induction of specific, associative behavioral memory. *Neurobiology of Learning and Memory* 80: 268–284.
- Xerri, Christian, Stéphanie Bourgeon and Jaques-Olivier Coq (2005). Perceptual context-dependent remodeling of the forepaw map in the SI cortex of rats trained on tactile discrimination. *Behavioural Brain Research* 162: 207–221.
- Yip, Moira (1998). Identity avoidance in phonology and morphology. In *Morphology and its Relations to Phonology and Syntax*, S. LaPointe, D. Brentari, and P. Farrell (eds.), 216–246. Stanford, CA: CSLI Publications.
- Zanone, Pier G. and J. A. Scott Kelso (1992). The evolution of behavioral attractors with learning, nonequilibrium phase transitions. *Journal of Experimental Psychology, Human Perception and Performance* 18: 403–421.
- (1997). The coordination dynamics of learning and transfer. Collective and component levels. *Journal of Experimental Psychology, Human Perception and Performance* 23: 1454–1480.