Category competition drives contrast maintenance within an exemplar-based production/perception loop.

Andrew WEDEL
Department of Linguistics
University of Arizona
Tucson, AZ, 85721
wedel@u.arizona.edu

Abstract

The evolution of competing lexical categories is simulated within a model in which lexical outputs are organized as sequences of articulatory gestures. When exemplar-based categories compete for assignment and storage of incoming exemplars in a production/storage loop, contrast between categories spontaneously emerges and remains stable, driven by the differences in storage consistency between more contrastive and less contrastive variants. Further, when lexical outputs are biased toward use of previously produced gestures, the set of exemplars in the lexicon evolve to be derived from a small set of contrastive units used in combination, despite the absence of direct selection for contrast at the sub-lexical level.

1   Introduction

The probability of accurate information transmission is dependent on the perceptibility of difference between differently signifying forms, that is, *contrast*. The possible mechanisms by which contrast arises and is preserved in lexical forms, on the other hand, have been less clear. Many grammatical theories of the last century assume that the language faculty is constituted to directly optimize contrast in some way, e.g., (Martinet 1955), (Flemming 1995), and much computational work also operates within the assumption that contrast between units of form is maintained through some kind of direct monitoring and manipulation of contrast, e.g., (Lindblom 1986), (de Boer 2000). In all of these approaches, contrast is a property of forms. Here, I will present evidence within an exemplar model of lexical production and perception (Goldinger 2000), (Pierrehumbert 2001) that the fact of a distinction between categories themselves, rather than the forms that instantiate them, can be indirectly responsible for driving contrast preservation through the statistics of assignment of forms to categories.

Within exemplar models of linguistic category structure, the act of categorizing a percept does not strip that percept of all non-contrastive detail, e.g., (Johnson 1997), (Pierrehumbert 2001). When we take into account evidence that production of an output of a lexical category may be based on details of previously perceived instances of that category (Goldinger 2000), we see that a production-perception feedback loop is closed, in which details in what is perceived can be subsequently reflected in the details of what is produced (Pierrehumbert 2001), (Oudeyer 2002). Whenever a system exhibits variation among elements, selection of variants over some criterion, and subsequent reproduction of selected elements, the system will evolve through natural selection on the basis of that criterion. Hence, any factors within the production-perception loop that bias the distribution of forms that are produced, the distribution of forms that are perceived, or the way that percepts are categorized, will result in evolution of category contents. Within the model presented here, lexical categories are populated by exemplars that have been previously categorized as correspondents of that category, and the output of a given category follows a distribution defined in part by the range of exemplars of that category, e.g., (Goldinger 2000), (Pierrehumbert 2001, 2002). Outputs are recognized as correspondents of a given category by comparison to exemplars already stored within that category (Pierrehumbert 2001); see also (Luce and Pisoni, 1998). Because outputs of a category can be re-stored as new exemplars within that category within a community of speakers, any asymmetries in either the form of outputs, or the likelihood of recognition and storage of those outputs, will result in a shift in the contents of that category over time (Pierrehumbert 2001), (Oudeyer 2002), and (Wedel 2004).

Here, I show simulation results suggesting that contrast between distinct form-meaning pairings can arise indirectly from asymmetries in the

consistency of categorization of more contrastive, versus less contrastive outputs[1]. Because more contrastive outputs make up a relatively greater proportion of the regularly stored exemplars in a given category than less contrastive outputs, they should have a proportionally greater influence on the evolution of that category. This asymmetry in the statistics of recognition and storage results in biased evolution of categories towards greater contrast.

1.1 Contrast preservation through categorization in morphology.

Within linguistics, the notion that contrast maintenance is an indirect effect of contrast's effect on a hearer/acquirer's categorization behavior has been suggested by Pierrehumbert (2002), and by Gregory Guy (1996) on the basis of corpus data on preservation of morphological contrasts. Guy notes that data from production corpora will always underestimate the true extent of speakers' failure to produce a given meaningful contrast. For example, if a transcriber perceives the utterance 'I cook the chicken', in the absence of additional information s/he is likely to simply transcribe it as such, even if the speaker actually intended the sentence to be in past tense, but elided the [-t] past tense marker. Guy notes that language acquirers are no different from transcribers, such that the perception data from which a language learner develops a grammar will be biased towards the more contrastive utterances in the production data set. This steady selection of more contrastive forms in the categorized utterance set upon which acquisition is based should result in a tendency for grammatical processes to emerge that appear to function to preserve contrast, when they in fact only act to reproduce the patterns in the data set that the acquirer perceives.

1.2 Contrast maintenance as a form of niche specialization.

This mechanism for category separation through competition for category members is formally parallel to a proposed mechanism of sympatric speciation first proposed by Darwin (1859, chap. 4) and further developed in recent theoretical research on the effects of resource competition on the distribution of phenotypes in a population (Kondrashov and Kondrashov 1999), (Dieckmann and Doebeli 1999 and references therein). In this model of sympatric speciation (speciation in the absence of geographical separation), phenotypic

divisions within a population and subsequent speciation can be driven by inequalities in the degree of competition experienced by individuals lying at different points on a distribution of phenotypes relating to resource exploitation. Individuals exhibiting intermediate phenotypes compete against a larger fraction of the population, while more extreme phenotypes have fewer competitors, and therefore greater individual access to resources. The higher fitness of individuals lying at the extremes of a phenotypic distribution can eventually produce a split in the population along this phenotypic dimension, setting the stage for subsequent speciation.

The same statistical influence of resource competition on fitness has also been proposed to drive 'niche specialization' among separate species occupying overlapping niches (Schoener 1974), (Dieckmann and Doebeli 1999). For example, if two species that utilize an overlapping set of resources jointly colonize a new environment, they tend to evolve to specialize on different portions of the resource distribution. This is proposed to occur because phenotypic variants of each species that happen to focus on an extreme of the resource distribution experience less competition than those who prefer the center of the distribution.

Within the exemplar based model proposed here for contrast maintenance, lexical categories are formally parallel to competing species undergoing selection for niche specialization. A category will be less often matched with a percept that is also close to another category than a percept that is close to no other category. Further, because the matching behavior of a category is determined by its contents, a category will evolve to be more specific for those percepts most often identified as members of that category. In this way, categories will tend to evolve to split the available percept space evenly, minimizing regions of overlap (see (Pierrehumbert 2002) for additional discussion of overlap minimization in evolving exemplar-based categories).

2 Modeling contrast maintenance through category competition.

To provide a simple illustration of the phenomenon of category boundary maintenance through patterns of category assignment, I show below results of a simple simulation of two interacting categories. Each category contains ten numerical 'exemplars' that can vary in value within a range from zero to ten. In each round, each category produces all of its exemplars one by one. Noise is introduced in production by adding to each output a random value between +/-0.4; increasing the amount of noise in outputs increases

---

[1] The simulations presented here assume that the result of lexical access is a unique output-category match, e.g., (Luce and Pisoni 1998).

the distribution of exemplars within a category, and increases the rate of change in category contents over simulation cycles.

Produced exemplars are subsequently re-stored in one of the two categories, based on how close the produced exemplar is to the average of each category. For exemplars lying between the averages of the two categories, the likelihood of assignment to a given category is proportional to the relative distance to that category's average. For example, if the average exemplar values in the two categories are 5 and 7 respectively, a produced exemplar with the value 3 will be re-stored in the first category, while an exemplar with the value 6 has an equal chance of being stored in either category. In this way, the two categories can be said to compete for produced exemplars on the basis of their own contents. Re-storage is accompanied by random deletion of a previously stored exemplar, such that the number of exemplars in each category remains constant.

A simulation begins with each category pre-seeded with 10 exemplars, all with the value 5. Figure 1A below shows the evolution of the value averages of the two categories over 2000 rounds of production and storage. Note that the averages diverge immediately from their originally shared value of 5, and over the run of the simulation, occasionally approach one another, but never cross.

Figure 1A

Evolution under Inter-category Competition
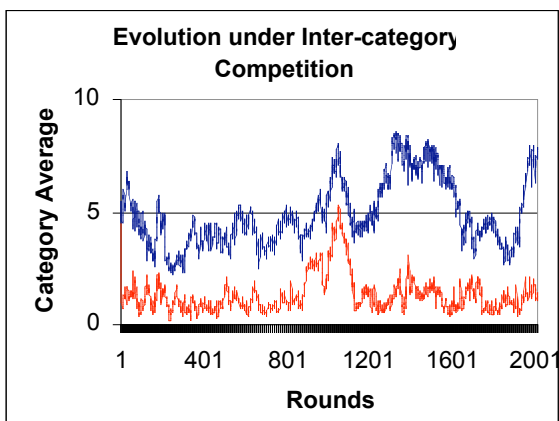
As a control, Figure 1B shows the results of a similar simulation in which categories do not compete for exemplars, but where exemplars are always re-stored in their category of origin. In this case, the category averages approach and cross one another multiple times, as we expect, given that their pathways through the simulation are independent.

Figure 1B.

Evolution with No Inter-category Competition
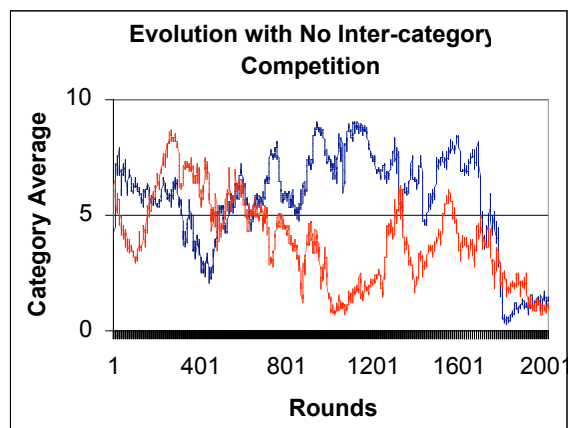
This difference is robust: when a simulation like that shown in Figure 1A was run 100,000 rounds, there were no crossovers of category averages, while 100,000 rounds of a simulation like that shown in Figure 1B produced 491 crossovers. The failure of the averages to cross in simulations in which categories compete for exemplars is due to the fact that exemplars located between the category averages are less often stored in any given category than exemplars lying to one side. Because the average value of categories depend on what has been previously stored, categories tend to shift their averages away from each other over time, in effect minimizing competition between them.

These simulations can be instructively compared to Janet Pierrehumbert's simulations of merger between exemplar-based phonetic categories under a leniting bias in a similar production/storage loop (2001, 2002). In Pierrehumbert's simulations, the number of outputs of a category in a given round is proportional to the number of exemplars stored there, such that a category that is more successful in competing for outputs grows more highly populated at the expense of another that is less successful, with the result that it may eventually subsume the less successful category, which then ceases to exist. If we assume that the notion of contrast only has functional substance in the context of an actual form-meaning pairing, then this is reasonable, because sub-morphemic categories do not have independent meaning of their own, but only contribute to marking a meaning difference in larger sound sequences. In contrast, in the simulations detailed above, the two categories always produce the same number of outputs in every round, regardless of their historical competitive success. This behavior seems reasonable for lexical categories, which can be anchored in physical and social experience outside the linguistic system. As such, distinctions in meanings may remain intact, even if

corresponding forms merge. For example, if the lexical form for the category 'sand' were to merge with that for 'hand', we presumably would not suddenly find ourselves significantly less interested in talking about 'sand' than we were before. The same cannot be said for the merger of sub-morphemic categories, as they have no independent correspondent in meaning.

3 The interaction of contrast maintenance and motor consolidation.

Given the hypothesis that contrast maintenance is driven through category competition between form-meaning pairings, we then need an account for the observation that morphemes themselves do not appear to be the minimal unit of contrast in phonological systems. Rather, we find that phonological systems can be described in terms of sub-morphemic contrastive features and featural groupings. In the simulations presented in this section, I explore the possibility that contrastive, sub-morphemic units can arise indirectly through category competition for the larger lexical forms of which they are a part.

The simulation architecture employed here consists of a single speaker/hearer pair, each equipped with small lexicons of categories populated with stored exemplars. In a given round, one of the pair utters the contents of its lexicon to the other, which attempts to categorize and store each utterance by comparing it to the exemplars stored in its lexicon. The sophistication of the simulation architecture is purposely kept very low to enable us to better assess the hypothesis that the simulation results are due to patterns of information flow in the system, rather than particular details of implementation.

The simulation architecture used here is based on the proposal that the articulatory gesture is a basic unit of phonological organization (Browman and Goldstein 1990), (see also Oudeyer 2002). Further, these simulations model the finding that practiced gestural targets serve as attractors in motor planning and execution, e.g., (Saltzman and Munhall 1989), (Shadmehr and Bashers-Krug 1997), and (Bybee 2001). Stated another way, because motor behavior is reinforced by repetition, the more highly practiced a given movement pattern, the more likely a future movement will follow that pattern. The neural mechanism for the consolidation of motor patterns is not under study in the simulations below, and so for computational simplicity it is simply stipulated. Simulations that model neural mechanisms for the development and action of attractors in developing systems can be

found in e.g., (Guenther and Gjaja 1996) and (Oudeyer 2002).

At the start of a simulation, each speaker/hearer is provided with a starting lexicon of comprising a number of categories, each containing 9 exemplars. Exemplars are structured as an ordered sequence of gestural targets. Two articulators are provided in the simulation, labeled X and Y, where X can vary in a range from 0.00 – 0.30, and Y from 0.00 – 0.10. Each exemplar consists of four ordered pairs ('segments') of articulator targets, as for example:

[X: 0.03] [Y: 0.10]; [X: 0.21] [Y: 0.03];
    [X: 0.12] [Y: 0.06]; [X: 0.00] [Y: 0.08].

A tight linkage between acoustic/perceptual and articulatory maps is assumed (e.g., (Oudeyer 2002) and references therein), and because structure potentially emerging from the interactions between perceptual and articulatory mappings is not at issue here, recorded exemplars and outputs of production are both encoded in the same units for computational simplicity.

Unless otherwise specified, at the beginning of each simulation the lexicons of each speaker/hearer are seeded with fully randomized exemplars. In a given round, one of the pair produces three randomly chosen exemplars from each of its lexical categories for the other, which categorizes and stores the produced outputs by comparison to the stored contents of its own lexical categories.

Production proceeds by selecting a single random exemplar from a category, and then assembling a corresponding output. To simulate the warping of motor targets toward more highly practiced outputs, each speaker/hearer retains a record of what articulatory targets have been produced over the previous six rounds. An output target value for each target value recorded in the chosen exemplar is established by comparing the reference exemplar target value to every target value recently produced by that articulator. Recently produced target values are activated in Gaussian proportion to their proximity to the reference exemplar target value[2].The probability that a particular target value will be chosen is

---

[2] The formula used to calculate the activation of a stored target value through its frequency of previous use and proximity to the corresponding exemplar value is:

$$A = n(2.141^{(-25(a-b))^2})$$

where n is the number of times that the target value under consideration has been produced in the previous six rounds, a is the target value, and b is the reference target value in the exemplar under current production.

directly proportional to its activation with respect to the corresponding position in the exemplar under production.

For example, if the particular target value recorded in the exemplar chosen as the basis for production has been produced often by the current speaker, it is likely to be faithfully reproduced in the output. On the other hand, if it has rarely been produced, but a nearby target value has been produced more often, the actual output target will likely match the more commonly produced, nearby target. The result is a steady tendency to consolidate motor patterns over time.

Finally, each production target is produced with Gaussian noise: in the simulations shown here, the variance in the noise distribution was such that each intended target had a 10% chance of being modified +/-.01 on the target scale. Up to a point, increases in the breadth and amplitude of the noise distribution increase the rate at which the system explores new states; beyond this point, the system begins to lose stability as production events become increasingly random.

Category assignment on the part of the hearer proceeds by comparing the speaker output to all exemplars stored in the hearer's lexicon. Whether a hearer exemplar will be counted as matching the speaker's output is assessed by comparing each output target value to the corresponding target value in the stored exemplar, where the probability of target matching follows a normal distribution with a match probability of 1 at equal target values and a standard deviation of 0.1 on the scale of possible target values.

After all matches have been determined within the hearer's lexicon, the output is assigned to and stored in a single lexical category, where the probability of category assignment is proportional to the square of the number of matching exemplars in each category[3]. For example, if an output is successfully matched to 2 exemplars in lexical category A, and 1 exemplar in B, it is four times as likely to be assigned to A as B. This matching and assignment procedure is intended to approximate probabilistic activation and competition between lexical entries in lexical access (Luce and Pisoni, 1998). If an output lies an equivalent distance between two categories (i.e., it is matched to equal numbers of exemplars in the two categories), it has an equal chance of assignment to each category. Only if a speaker output is matched to *no* exemplar in the hearer's lexicon will it fail to be assigned to any category. When an output is assigned to a category and stored there as a new exemplar, a randomly chosen exemplar from a previous round is discarded.

3.1 Evolution of lexical categories in the absence of inter-category competition.

The tendency to warp output target values toward those values that have been produced before results in a steady reversion to mean target values over the course of the simulation, counteracting the dispersive effects of noise in production (Wedel 2004)[4]. This can be seen in a simulation in which category competition in lexical access is disabled by providing the speaker/hearers with an additional channel of communication, such that each speaker output is stored directly in the corresponding hearer category without regard to similarity any other category's contents. Figure 2A shows the distribution of values for the X target over four lexical categories from one of the speaker/hearers at the beginning of such a simulation, and 2B the distribution of X target values over these lexical categories after 1000 rounds.

While the target values for both articulators are distributed across their possible ranges at the beginning of the simulation, steady feedback pressure for output target values to become more alike results in the evolution of a system with only one possible target value for each articulator, and as a result, all lexical categories evolve to contain the same set of exemplars. Figure 2C shows the consensus target values for each of the four lexical categories at round 1000. At 1000 rounds, Y values have settled around a single value as well (not shown).

---

[3] Varying the scaling between relative number of matches and probability of category assignment within a reasonable range does not change the behavior of the simulation. As the exponent is raised, variation in category assignment decreases, resulting in slower evolution of the system. As the exponent is lowered on the other hand, category assignment becomes less dependent on the relative goodness of match.

[4] Reversion to the mean in production can also be created by selecting multiple exemplars from a category and averaging them in production (e.g., Pierrehumbert (2001)). The category-sharpening effect of motor consolidation renders this unnecessary in this model.
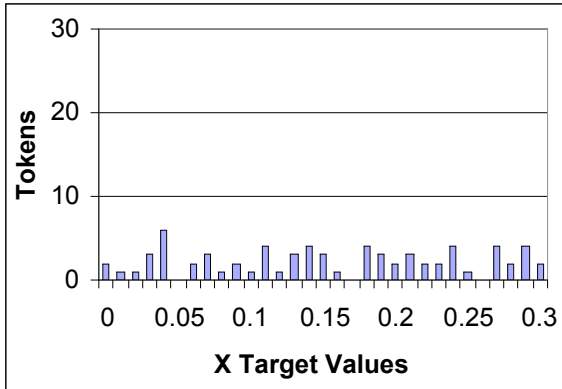
Figure 2A. X values at Round 0
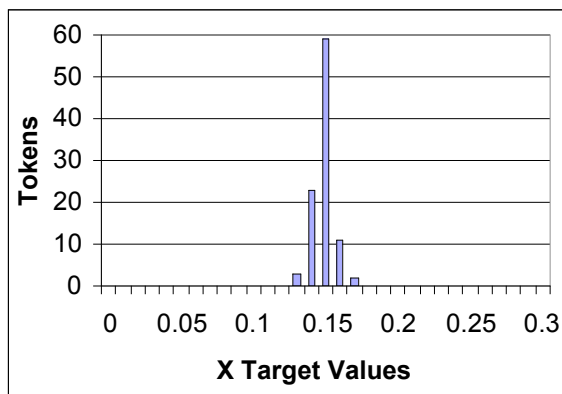

Figure 3A. X values at Round 1000.

Figure 2B. X values at Round 1000
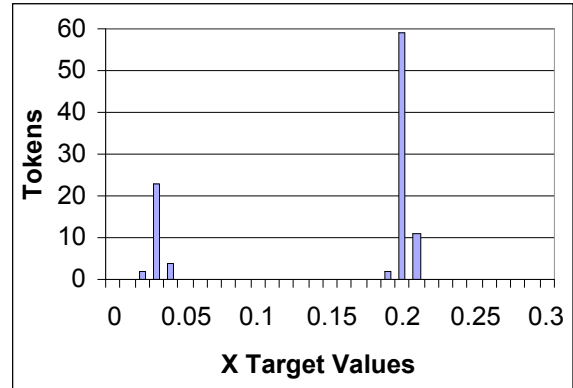

Figure 3B. Consensus target values: Round 1000.

Category	X/Y	X/Y	X/Y	X/Y
A	.20/.06	.20/.05	.20/.05	.03/.06
B	.03/.06	.20/.06	.20/.06	.20/.06
C	.03/.05	.03/.06	.20/.06	.21/.06
D	.20/.06	.20/.06	.20/.06	.20/.06

In Figure 3B, equivalent 'segments' are shaded equivalently. Each lexical category is distinct from every other. In runs with larger numbers of lexical categories, two Y target values often develop as well, assorting with two or more X target values to provide sufficient numbers of contrastive units (not shown).

Note that this effect is not dependent on less contrastive variants being less often stored. In the numerical simulations described above in section 2, every output was stored in a category. In the simulations described in this section, every variant form was stored in some category, provided it could be matched to at least one exemplar. A variant lying between two categories has then in fact a greater chance of being matched and stored than a variant that lies an equivalent distance away from a single category, even though the latter is functionally more contrastive. This generous assignment and storage procedure was chosen to make it less likely that the development and maintenance of contrast within the simulations could be due to differential *rates* of storage, as opposed to differential *consistency* of storage. Previous work (Wedel 2004) showed that in similar simulations in which outputs matching multiple categories were at a disadvantage in storage efficiency relative to those matching just one category, development and maintenance of contrast was yet more robust than in the simulations shown here. Phenomena such as the neighborhood density effect in lexical access indicate that outputs activating multiple lexical categories may in fact be at a disadvantage in

Figure 2C. Consensus target values: Round 1000.

Category	X/Y	X/Y	X/Y	X/Y
A	.14/.02	.14/.02	.14/.02	.14/.02
B	.15/.02	.14/.02	.14/.02	.14/.02
C	.14/.02	.14/.02	.14/.03	.13/.02
D	.14/.03	.14/.02	.14/.02	.14/.02

3.2 Inter-category competition supports maintenance of contrast.

When category competition is reintroduced, however, something quite different happens: although target values still show significant consolidation over the course of the simulation, sufficient distinctions remain to preserve contrast between lexical categories. Figure 3A shows the range of X target values, and 3B the lexicon at round 1000 of a simulation that incorporates category competition..

recognition (reviewed in (Luce and Pisoni 1998)). However, the simulation results presented here suggest that in addition to the effect of any disadvantage in storage of poorly contrastive forms, lower storage consistency of less contrastive forms alone can contribute to contrast maintenance between lexical categories.

3.3 Local lexical contrasts can support global sub-morphemic contrasts.

Because similar target values are bound together in a single unit of motor production, a given target value that is required to maintain contrast between two categories can support the persistence of that target value elsewhere, even if it would otherwise tend to merge with another nearby target value. To illustrate this phenomenon, Figures 4A and B compare the results of simulations beginning from two distinct, non-random starting points. Each simulation uses eight lexical categories, pre-seeded with exemplars consisting of three X target values ([0.30], [0.15], and [0.05]), and one Y value [0.00]. The X [0.05] value (found only in categories A, B, and C) is in the minority relative to the [0.15] target value, and therefore will be under pressure from motor consolidation to merge with it.

In the starting lexicon shown in Figure 4A, the minority X value [0.05] can increase and merge with the more frequent X value [0.15] without any loss of category contrast, because categories A, B and C differ from every other category by at least two positions. (The identical Y values across the lexicons in Figures 4A and B are irrelevant to contrast, and so are omitted for clarity).

Figure 4A. Starting target values: Round 0.

Category	X	X	X	X
A	.30	.15	.30	.05
B	.15	.15	.30	.05
C	.30	.30	.30	.05
D	.15	.30	.15	.30
E	.15	.30	.30	.15
F	.15	.30	.30	.30
G	.30	.30	.15	.15
H	.30	.30	.30	.30

For example, if the X target value [0.05] in category A were to increase to [0.15], category A would still remain distinct from every other category. When simulations are run with this starting lexicon, the minority [0.05] values do in fact tend to quickly merge with the more frequent [0.15] value. In ten runs of the simulation with this starting lexicon, merger always took place within 400 rounds.

Figure 4B shows a starting lexicon that is nearly identical, except that if the [0.05] target values of categories A and B increase to [0.15], these categories will lose contrast with categories D and E, respectively. The [0.05] target value of category C can still merge with [0.15] with no loss of category contrast.

Figure 4B. Starting target values: Round 0.

Category	X	X	X	X
A	.15	.15	.30	.05
B	.30	.15	.30	.05
C	.30	.30	.30	.05
D	.15	.15	.30	.15
E	.30	.15	.30	.15
F	.30	.30	.30	.30
G	.30	.30	.15	.15
H	.15	.30	.30	.30

In this case, in ten simulations run out to 1000 rounds, the [0.05] target values of categories A and B did not ever merge with the [0.15] target value, as expected. Interestingly, in eight of these runs, the [0.05] target value of category C also failed to merge, even though this value was not required for contrast in category C. Note that the starting lexicons shown in Figures 4A and B have the same relative numbers of [0.30], [0.15] and [0.05] X target values; it is only their distribution in the lexicon, and therefore their functional load that is different. The difference in evolutionary pathways of these lexicons illustrates that within this simulation architecture, category competition resting on a given target value in one part of the lexicon can stabilize that target value throughout the lexicon, even in regions of the lexicon where its functional load is low.

As suggested by the simulation architecture, this effect is dependent on the details of frequency: *ceteris paribus*, the ratio of contrast-bearing to non-contrast bearing instances of a target value influences the probability of merger (see (Labov 1994:328ff) for similar arguments on mergers in vowel systems). In the lexicon in 4B for example, two of the three instances of the target value [0.05] functioned to maintain lexical contrast. If the lexicon is altered such that only *one* of the three instances of that target value bears responsibility for contrast in a lexical category, then in a significant number of runs, the other two instances of the [0.05] value merge with [0.15], and in the process drag the contrast-bearing [0.05] value with them, producing a pair of 'homophonic' categories in the lexicon with identical contents (not shown). Within this simulation architecture, the attractor formed by identical target values in output assembly is stronger than the statistical force pressuring different category contents to diverge in storage, such that once homophonic categories

form, they never split and regain contrast. This appears to be largely true for actual lexical categories as well (but see e.g., (Yaeger-Dror 1996) and (Jurafsky et al. 1996) for evidence that homophonous categories may be able to split under some circumstances).

4   Discussion

The simulation results described here illustrate that competition between categories for form variants in a production/storage loop indirectly supports maintenance of contrast across form-category pairs. As suggested by (Guy 1996) and (Pierrehumbert 2002) in linguistics, and by theoretical work in niche specialization in evolutionary biology (Schoener 1974), (Deickmann and Doebeli 1997), this phenomenon rests on unequal partitioning of variants across self-reproducing categories: those variants that are split among multiple categories contribute less to the evolving form of any given category than those variants that are more consistently stored, with the result that the lexicon evolves to reflect the more contrastive variants.

Is this a potential mechanism contributing to contrast maintenance at the sub-morphemic level? In support of this possibility, speakers have been shown to produce more contrastive phonetic detail when producing words in high-density lexical neighborhoods, e.g., (Goldinger and Summers 1989), (Wright 1996), and (Brown 2002), which could be a reflection of the distribution of phonetic details stored in the lexical categories in high density neighborhoods.

The structure of exemplars in the simulations in section 3 reflect the fact that lexical forms consist of multiple, temporally ordered articulatory gestures. This structure allows, but does not dictate, the development of a combinatorial system in which simulated gestures or gestural groupings are reused in distinct lexical forms. However, evidence strongly suggests that practice of coordinated muscular gestures results in consolidation into larger-scale motor programs, which then serve as attractors in motor planning and execution (Shadmehr and Bashers-Krug 1997 and references therein). To simulate this effect, in output assembly target values and larger groupings of target values were warped towards those values that had been frequently produced in the speaker's recent history. The resulting tendency to minimize target value differences conflicts with the statistical reward enjoyed by more contrastive forms, resulting in a optimization in which lexical entries evolve to contain contrastive exemplars, which are themselves composed of a small number of contrastive units. Rather than being stipulated

anywhere in the system, the contrastiveness of these sub-lexical units evolves indirectly through competition between the actual form-category pairs that contain them. Interestingly, as we saw in the simulations described in conjunction with Figures 4A and B, the tight association of target values as 'motor units' allows a functionally unnecessary contrast to persist in a given lexical category, if that contrast is functionally required in another.

The potential influence of contrast on sound change suggested by these results is supported by at least two well-described phonological patterns. First, certain exceptions to otherwise regular sound change, sometimes referred to as 'anti-homophony effects', occur precisely where sound change would give rise to loss of a paradigmatic contrast. In this case, data from unrelated languages supports a cross-linguistic tendency for contrastive exemplars to be preferred exactly where lexical categories are in greatest competition (Blevins to appear), and (Gessner & Hansson 2004). A second finding is that rare phonological contrasts (e.g., a three-way contrast in vowel or consonant length, or a three-way contrast in nasalization) are not randomly distributed in the lexicon. Rather in languages making use of rare contrasts, these contrasts are frequently the exponents of contrastive morphological features, and hence are only contrastive in contexts of lexical competition (Blevins 2004, chapter 8). The simulations presented here can account for both these findings in terms of contrast-driven statistical selection of exemplars at category extremes. In the case of antigemination, this selection inhibits the progression of a sound change in limited contexts where morphemes compete. In the case of rare contrasts, selection inhibits expected mergers in limited contexts where that contrast is the sole exponent of contrast.

Finally, the results of these simulations contribute to an ongoing discussion of the divergent relationship between lexical frequency and morphological, versus phonological 'regularity' (see e.g., Pierrehumbert 2002). The well-known tendency to morphological irregularity in high-frequency forms can be explained as an effect of frequency on lexical access: the higher resting activation level of frequent forms should allow them to be identified holistically, rather than through identification of their individually contrastive morphemes. Similarly then, we might expect that highly frequent words should be able to evolve to be phonologically exceptional, for example by resisting a sound change sweeping through the rest of the lexicon, or developing an otherwise unattested phone. However, in general

we find just the opposite: highly frequent forms do conform to sound changes initiated elsewhere, tend to be the most lenited, and tend as well to comprise more common sounds (see Bybee (2002) and Pierrehumbert (2002) for discussion). This can be explained within a model in which phonological, but not morphological categories tend to be coextensive with motor units, if we assume a tendency toward effort minimization in production. *Ceteris paribus*, highly practiced motor scores are deployed more rapidly and accurately than less practiced motor scores (Shadmehr and Bashers-Krug 1997 and references therein), and so we can consider lexical evolution toward use of more common motor scores a form of lenition, defined with respect to the particular language use of speakers in a speech community. If a sound change sweeps through most of a lexicon altering a motor score into another, highly frequent forms might also be expected to shift away from the original, now infrequent motor score to the new highly frequent one, not because they must, but because their high frequency encourages lenition (reviewed in Bybee 2001). Preliminary simulation results exploring the interaction of exemplar frequency and warping toward frequent target values support this hypothesis.

5   Acknowledgements

References

J. Blevins. (to appear) Understanding antigemination. In *Papers from* the *International Symposium on Linguistic Diversity*. Zygmunt Frajzyngier ed., John Benjamins, Amsterdam.

J. Blevins. 2004. *Evolutionary Phonology: The emergence of sound patterns*. Cambridge University Press, Cambridge.

C. Browman and L. Goldstein. 1990. Gestural specification using dynamically-defined articulatory structures. *Journal of Phonetics* 18:299-320.

R. A. Brown. 2002. Effects of Lexical Confusability on the Production of Coarticulation. *UCLA Working Papers in Linguistics*, no. 101.

J. Bybee. 2001. *Phonology and language use*. Cambridge University Press, Cambridge.

J. Bybee. 2002. Word frequency and context of use in the lexical diffusion of phonetically conditioned sound change. *Language Variation and Change* 14:261-290.

C. Darwin. 1859. *On the origin of species by means of natural selection*. J. Jurray, London.

B. de Boer. 2001. *The origins of vowel systems*. Oxford University Press, Oxford.

U. Dieckmann. and M. Doebeli. 1999. Theoretical considerations of sympatric divergence. *American Naturalist* 107:256-274.

E. Flemming. 1995. *Auditory Representations in Phonology*. Ph.D. Dissertation, University of California, Los Angeles.

S. Gessner and G. Hansson. 2004. Anti-homophony effects in Dakelh (Carrier)Valence Morphology. To appear in *Proceedings of Berkeley Linguistics Society* 30.

S. D. Goldinger. 2000. The role of perceptual episodes in lexical processing. In *Proceedings of SWAP (Spoken Word Access Processes)*, A. Cutler, J. M McQueen, and R. Zondervan, ed., pp. 155-8. Max-Planck-Institute for Psycholinguistics, Nijmegen.

S. D.Goldinger and W. V. Summers. 1989. Lexical neighborhoods in speech production: a first report. *Research on Speech Perception Progress Report*, No. 15:331-342. Bloomington.

F. H. Guenther and M. 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.

G. R. Guy. 1996. Form and function in linguistic variation. In *Variation, Change and Phonological Theory*. G. Guy, C. Feagin, J. Baugh, D. Schiffrin, and M. B. Kac., ed., pp. 125-43. Benjamins, Amsterdam.

K. Johnson. 1997. Speech perception without speaker normalization. In K. Johnson. and J. W. Mullennix, ed., *Talker Variability in Speech Processing*. Academic Press, San Diego.

A. S. Kondrashov and F. A. Kondrashov, 1999. Interactions among quantitative traits in the course of sympatric speciation. *Nature* 400:351-354.

W. Labov. 1994. *Principles of linguistic change, Vol. 1, Internal factors*. MA: Blackwell, Oxford and Cambridge.

B. Lindblom. 1986. Phonetic Universals in vowel systems. In *Experimental Phonology*. J. J. Ohala and J. J. Jaeger, ed., pp. 13-44. Academic Press, Orlando.

P. A. Luce and D. B. Pisoni. 1998. Recognizing Spoken Words: the neighborhood activation model. *Ear and Hearing* 19:1-36.

A. Martinet. 1955. *Economie des changements phonétiques*. Francke, Berne.

P-Y. Oudeyer. 2002. A Unified Model for the Origins of Phonemically Coded Syllable Systems. In *Proceedings of the 24th Annual Conference of the Cognitive Science Society*. B. Bel and I. Marlien, ed., Laurence Erlbaum Associates.

J. Pierrehumbert. 2001. Exemplar dynamics: Word frequency, lenition, and contrast. In *Frequency effects and the emergence of linguistic structure*. Bybee, J and P. Hopper, ed. pp. 137-157. John Benjamins, Amsterdam.

J. Pierrehumbert. 2002. Word-specific phonetics. In *Laboratory Phonology 7*. C. Gussenhoven and N. Warner, ed. Mouton, Berlin; New York.

E. Saltzman. and K. G. Munhall. 1989. A dynamical approach to gestural patterning in speech production. *Ecological Psychology* 1:333-382.

R. Shadmeh. and T. Bashers-Krug. 1997. Functional stages in the formation of human long-term motor memory. *Journal of Neuroscience* 17: 409-419.

T. W. Schoener. 1974. Resource partitioning in ecological communities. *Science* 185:27-37.

A. Wedel. 2004. *Self organization and categorical behavior in phonology*. Ph.D. dissertation, University of California, Santa Cruz.

R. Wright. 1996. Lexical Competition and Reduction in Speech: A Preliminary Report. In *Research On Spoken Language Processing: Progress Report* No. 21 (1996-1997), Indiana University.

M. Yaeger-Dror. 1996. Phonetic evidence or the evolution of lexical classes: The case of Montreal French vowel shift. In *Towards a Social Science of Language*, G. Guy, C. Feagin, J. Baugh, and D. Schiffrin, ed. pp. 263-87. Benjamins, Amsterdam.