

Self-Organization in Phonology

Andrew Wedel, University of Arizona

Structure can arise in a system in many different ways. Self-organization is one general mechanism for structure formation which has relatively recently been explored as a possible contributor to patterns found in language. The aims of this chapter are:

- to provide an overview of self-organization as a general mechanism for structure formation;
- to describe some of the ways that self-organizational processes can interact with other familiar mechanisms for structure formation;
- to review selected work done to date arguing that particular phonological phenomena may arise through the contribution of self-organizational mechanisms.

Self-organization has been argued to play a role in a wide variety of phonological phenomena, from the development of a phonological grammar in acquisition through the systemic grammatical changes that occur in language over many generations. A central mechanism in self-organization is feedback, in which the properties of a current state of a system are dependent in some way on those of a preceding state. When such states are linked in a temporal chain, structure can develop in ways that cannot be described in terms of a single set of causal steps. Arguments that phonological patterns can emerge in this way are not new: Lindblom et al. (1984) for example argued that phonological inventories are shaped by interactions between phonological categories, mediated by constraints on effort and perceptibility rather than through, for example, innate constraints on features and their combinations. In the last decade, however, the scope of self-organization based accounts has rapidly widened to include phonological acquisition, evolution of complex grammatical patterns, evolution and maintenance of phonological contrast, grammatical ‘conspiracies’ and many other phenomena. A representative set of these accounts will be discussed in detail below in section 3, but it will be helpful to begin with a general introduction to key concepts in self-organization as a general structure formation mechanism. Supplementary simulations illustrating these concepts are included in the online version of this article.

1. Self-organization as a pathway to structure

‘Self-organization’ is not a concept with crisp edges. Instead, it is a big-tent term covering the many ways structure can form in *non-linear, dissipative* systems. Non-linear systems are those in which the properties of the system as a whole cannot be understood in terms of the properties of individual system elements, in other words, in which new properties emerge through interaction. Cooked egg-white is a familiar example of a non-linear system. Egg-white consists primarily of the protein albumin, and prior to cooking the globular albumin molecules slide past one another easily, producing a translucent semi-liquid. When the temperature is raised beyond a certain point, the albumin protein chains unfold and stick to each other, creating a large, highly interlocked structure that is opaque and stiff. These properties of cooked egg-white cannot be understood in terms of the summed properties of individual unfolded albumin molecules, but from their interaction as they stick together. Likewise, many features of language are dependent on the interaction of elements at some lower level of description for their existence, most obviously in that language change arises through language use between individuals and/or transmission across generations.

Dissipative systems are those in which a given state or structure is maintained through a constant flux of energy or matter. As a consequence, an account for a structure within a dissipative system includes time at some level. A ripple in a creek provides a familiar example of a higher order structure produced through a flux at a lower level of description. At one level of description, a ripple is an independent element of a creek that can interact with other elements at that level, such as a floating leaf or another ripple. At a lower level of description, it is a vast and constantly changing set of water molecules interacting with each other and the creek bed as they move. If the flow of water stops, the ripple disappears. Any given language is

dissipative in the sense that it is instantiated through usage and transmission, just as a ripple is instantiated through flow. As we will see in the examples below in section 3, self-organizational accounts of structure formation in language depend on cycles of use and/or acquisition.

Structure arises in non-linear, dissipative systems when many similar elements or events interact over time to produce persistent changes at some higher level of organization. Typically, structure formation in self-organizing systems is the result of positive and negative feedback loops engendered by the interaction of system elements with each other or with the environment. Positive feedback (also sometimes referred to as autocatalysis) arises when a given event makes a similar event more likely in the future. An example is the population growth that occurs when individuals have offspring at greater than the replacement rate. In this case, the birth of each additional individual makes a subsequent birth more likely. Positive feedback promotes change and can result in runaway processes. As an example of a potentially central role of positive feedback in language, I have argued that positive feedback in the form of similarity bias in production and perception may drive the development of coherent grammatical patterns despite storage of low level phonotactic detail in the lexicon, which should otherwise promote lexical idiosyncrasy (Wedel 2007).

Negative feedback arises when an event makes a similar event less likely in the future, as when a growing population outstrips its supply of resources. In this case, each additional birth *lowers* the probability of a subsequent birth through increased competition. Negative feedback promotes stability. Both positive and negative feedback represent types of non-linearity, because the description of patterns resulting from feedback must make reference to interactions between system elements. Self-organization often occurs in systems through positive feedback between system internal elements that is prevented from snowballing beyond a certain point by negative feedback. Examples of this sort include population growth limited by finite resources, thunderstorm structure in which a growing updraft is constrained by a resulting downdraft, and economic bubbles, burst by the collapse of credit. In section 3, we will see several examples of the ways negative feedback may inhibit loss of phonemic contrast over the course of language change.

Self-organized systems frequently exhibit *emergence*. Emergence in this context refers to the generation of a higher order structure that interacts meaningfully with other structures of the system at this level of description. Our earlier example of a ripple in a stream serves as a familiar instance of emergence: the influence of a ripple on other system elements (a leaf, another ripple) is most usefully described in terms of our understanding of the behavior of ripples, rather than our understanding of water molecules. Likewise in language, for example, many phonological patterns can be described in terms of the interaction of (possibly conflicting) phonotactic generalizations or constraints. In models in which phonotactic generalizations are derived over the course of acquisition and usage from the lexical items that instantiate them, these generalizations are emergent (see e.g. Blevins 2004, Wedel 2006, Mielke 2008).

Finally, self-organized systems frequently exhibit phase transitions between semi-stable states defined by *attractors*. An attractor is a system state (or set of states) that nearby states tend to evolve toward. A simple visual metaphor for a system with multiple attractors is a surface with multiple basins. If a ball is placed somewhere on this surface, it will tend to roll to the bottom of whatever basin it happens to be in. Phase transitions correspond to the transition from one basin of attraction to another and are accompanied by a shift in the behavior of a system. In our visual analogy, if we begin to shake the surface, the ball will begin to bounce around within its basin and may eventually by chance roll up and over into a new basin, where it remains until it rolls up and over into a new basin. Within the domain of morphology, local similarity bias in the form of analogical extension has been argued to create attractors that influence the course of morphological change over time (see e.g. Hock 2003, Garrett 2008). Likewise, pockets of formally similar irregulars ('gangs') have been shown to be more likely to recruit new members than formally isolated irregulars (Bybee & Moder 1983, Stemberger & MacWhinney 1988). Under this general model, coherent generalizations over forms act as emergent attractors. These patterns of similarity-based extension and, plausibly, resistance to extension are consistent with a model in which local similarity effects play a significant role in the formation of larger-scale morphological regularities. For a detailed implementation of this type of model simulating the evolution of past tense forms in Old English, see Hare & Elman (1995). In a similar fashion, I have argued that similarity biases at the level of sound-categories

may underlie the development of regular patterns in phonology (Wedel 2007) as well as the outcome of conflicts between phonological and morphological regularities (Wedel 2009).

The Game of Life (J.H Conway, reported in Gardner 1970) provides a simple example of a deterministic, self-organized system that exhibits all these properties. (Examples of The Game of Life can be found on the web; see also the online version of this article). The Game of Life is a simple cellular-automaton that occupies an infinite, two-dimensional orthogonal grid, the cells of which can either be 'alive' or 'dead'. There are three simple rules governing cell-birth and death, each of which makes reference to a cell's eight immediate neighbors:

1. If a living cell has fewer than two living neighbors, it dies.
2. If a living cell has more than three living neighbors, it dies.
3. If a dead cell has exactly three living neighbors, it becomes alive.

The grid is initialized with some seed pattern of living cells, and then left to evolve according to these three rules. Some seed patterns result in uninteresting outcomes: if the distribution of living cells is too sparse, all cells quickly die; conversely, some seed patterns are stable and do not change even though the rules continue to be applied. (Four cells arranged in a square is one example of such a stationary pattern.) Other seed patterns produce oscillating structures or structures that move in a consistent direction across the grid. 'Gosper's Glider Gun' is a particularly beautiful example of the complexity that can arise through the interaction of these simple rules over time. (You can find a movie of Gosper's Glider Gun on the web or in the online version of this article).

The Game of Life exhibits many of the typical features of self-organizing systems. Structure-formation depends on the interaction between elements (it is non-linear) and the application of cell birth- and death-rules over time (it is dissipative). It also requires the interaction of context-dependent positive and negative feedback processes: depending on the local context, the birth of a cell can cause the birth of additional cells in the next round or it can cause death. As in many self-organizing systems, structure arises in The Game of Life through positive feedback that is held in check by negative feedback. Finally, this system exhibits emergence, in which distinct groupings of living cells function as units with predictable behavior. In Gosper's Glider Gun for example, two large groupings of cells bounce off of stationary square groupings at the edges of the system, and then bounce off of each other. In the process of bouncing off of each other, they create small self-contained 'gliders' that embark on an infinite journey away from the center.

2. Self-organization in interaction with other influences on structure

Self-organization does not operate in a vacuum. It contributes structure in a context supplied by the system and its environment, and the properties of the system that support and direct the emergence of new structure can have any source, including innate pre-specification. For example, features of the environment can supply negative feedback or serve as templates that give initial direction to self-organized structure. In the Game of Life, self-organized structure formation is dependent on the properties of the environment (the orthogonal grid), and on the pre-determined character of interactions between cells. The structures that develop are also critically dependent on the initial seed pattern of living and dead cells which serves as an organizing template.

From a design point of view, self-organization is a powerful tool: if the details of a complex structure can be constructed through emergence instead of direct specification by some other means, it can be encoded considerably more compactly than otherwise possible (Gell-Mann 1992). For example, the specification of Gosper's Glider Gun requires only a description of the environment, the rules for cell birth and death, and the initial seed pattern. Further, different complex structures can be created by minimal changes to this description, such as the properties of the environment, the seed pattern or the rules for cell-birth and death.

Similarly, many biological structures are thought to emerge from self-organizational pathways which are given shape and direction by innately specified contexts. For example, the spots and stripes that are found

throughout the animal kingdom have been proposed to emerge through a single basic system with slight variations involving the competition between diffusing activator and inhibitor molecules in an animal's skin. (This model was originally proposed by Alan Turing in 1950; for discussion see Ball (1999), Chapter 4). Different shapes and patterns of coat markings can be produced simply by changing the relative diffusion speeds of the activator and inhibitor, or by changing the shapes of the underlying pigment-producing cells. This is an informationally much more compact way to produce a complex coat pattern than, for example, specifying the state of each and every individual pigment-producing cell. Furthermore, patterns that arise through self-organizational pathways are often very robust to perturbation because of the role that attractors play in the evolution of the system. For an overview of some of the many biological patterns that are thought to arise through self organization, including examples of the interaction between self-organizational pathways and other mechanisms, see Camazine et al. 2001.

There are many resources available on the web to learn more about self-organization and related concepts. Excellent published resources include *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity* (Kauffman 1995), *A Self-Made Tapestry* (Ball, 1999) and *Self-organization in biological systems* (Camazine et al. 2001).

3. Self-organization as a pathway for structure formation in language

Much recent work approaches language as a complex adaptive system in which grammatical patterns are emergent properties resulting from the repeated interaction of the many different elements that make up a larger language system: innate and acquired biases; forms at multiple levels of representation; interacting spheres of use; sociolinguistic networks; and chains of acquisition and transmission over longer time-scales. For representative examples, see Haspelmath (1999); Nettle (1999a, b); Plaut & Kello (1999); Browman & Goldstein (2000); de Boer (2000); Croft (2000); Lindblom (2000); Bybee (2001); Kirby (2001); Oudeyer (2002); Bod et al. (2003); Blevins (2004); Wedel (2007), Boersma & Hamann (2008), Mielke (2008); Blevins & Wedel (2009), Kirby et al. (2009).

A priori, there are at least two plausible reasons to think that self-organizational pathways may contribute to some of the patterns we find in language. The first is simply that at many levels and time-scales, language provides the necessary conditions to support spontaneous emergence of patterns through self-organizational pathways (see e.g. Lindblom et al. 1984; Ohala 1989; Lindblom 1992; Keller 1994; Labov 1994; Cziko 1995; Elman 1995, Dennet 1995, Deacon 1997, Cooper 1999; Hurford 1999; Steels 2000, Bybee 2001; Blevins 2004, MacWhinney 2006, Pisoni & Levi 2007, Beckner et al. 2009, and many others). Language involves the repeated interaction of many like elements, in like ways, at many levels of description and time-scales. Variation and bias in language acquisition and use provide many potential feedback pathways. Basic language-external constraints, from articulatory and perceptual factors through general categorization mechanisms to cross-culturally common salience relationships, all provide structures and templates that could give common shape to self-organized patterns. Because structure tends to spontaneously emerge under these conditions, it would be surprising if self-organizational pathways do not contribute to the formation of some of the many observed patterns in language, whether language-particular or crosslinguistically frequent. Put another way, if we found that self-organizational mechanisms played no role in the emergence of any observed language patterns, our burden would be to explain why not.

The second, perhaps less compelling reason derives from design-principles. As briefly reviewed above, the specification of a complex pattern can be much more compact and the resulting structure more robust to perturbation when created through self-organizational pathways. To the extent that the language faculty has evolved under functional constraints for use and that grammars continue to do so diachronically, self-organization represents a powerful mechanism for structure in the "blind watchmaker's toolbox" (Dawkins 1986).

Self-organization in phonology

Explorations of self-organizational accounts for linguistic patterns can be understood in terms of the general scientific goal of explaining more with less. Just as a good Optimality Theory (Prince & Smolensky 1993) account attempts to explain new patterns through rankings of existing constraints, an account for linguistic patterns making use of a self-organizational pathway attempts to explain a complex pattern through the interaction of simpler independent mechanisms. Often, authors of these accounts argue that a linguistic pattern emerges through the interaction of domain-general factors rather than through innate grammatical mechanisms, but it is worth emphasizing that this is not a necessary feature of a self-organizational account. Just as a self-organized biological pattern may arise from innately specified processes, a self-organized linguistic pattern could arise from simpler language-specific structures. Some models under the rubric of Biolinguistics have been framed in these terms; see e.g. Medeiros (2008).

Many self-organizational accounts make use of computational simulation, either as an existence-proof that a given structure can arise through interactions between some defined set of system properties, and/or as a supporting illustration for verbal or analytic arguments. Simulation is particularly useful in this context because self-organization proceeds through chains of circular causation progressively building structure over time. As a consequence, verbal descriptions of proposed self-organizational processes are often hard to assess critically. More importantly, interacting feedback loops are notorious for producing counter-intuitive results, so a computational implementation of a model provides both an important research tool for a theorist and a demonstration for a reader that the model operates as expected (Peck 2004). The following is a brief survey of several phenomena in phonology that have been proposed to arise in part through self-organizational pathways. This is of course just a sample of the many insightful self-organization-based models of phonological phenomena and interested readers are encouraged to explore the literature further.

Early phonological acquisition

In early phonological acquisition, initial relatively accurate word imitation is followed by a period of less accurate, but more systematic productions (Ferguson and Farwell 1975, reviewed in Vihman et al, 2009). This is reminiscent of the U-shaped learning curve of irregular morphological forms in which irregulars are initially reproduced faithfully, followed by a period of over-regularization, followed in turn by increasingly accurate production. Further, while these production patterns are consistent for a given child, they differ across children suggesting that the pathway to phonological competence is not prespecified at this level (Beckman & Edwards 2000, Vihman & Croft 2007, Vihman et al. 2009). Vihman et al. (2009) propose that this phenomenon can be explained in a model based on the ability of infants to acquire individual word-gestalts, combined with an ability to generalize over those gestalts through feedback from their own production (see also Pierrehumbert 2003 for related discussion). Under this model, an infant's set of practiced babbles provide the seed patterns for initial generalizations over learned word-gestalts. Feedback between these initial generalizations and word-productions allows the infant to develop practiced sublexical 'templates' (see also Fikkert 2007, Fikkert & Levelt to appear). The substitution of these generalized phonological 'templates' in place of gestalts accounts for the period of poorer production accuracy in matching the adult pronunciation of given words, yet greater precision between individual utterances. Accuracy subsequently improves as the myriad interactions with caregivers further shapes the trajectory of learning. Vihman et al. argue that a self-organizing, feedback-driven model of this kind is particularly well-suited to explain both the highly individual initial production templates observed in children, and their subsequent convergence on a community standard of pronunciation.

Conspiracies in historical phonology

Conspiracies, in which a seemingly disparate set of processes all result in a common pattern, are widespread in phonology (Kisseberth 1970). Robert Blust identifies many different types of diachronic change in Austronesian languages that conspire to create a disyllabic word, in many cases restoring a historical disyllable that had lost or gained a syllable through other changes (Blust 2007). In the history of Javanese, for example, reduplication, epenthesis, deletion and loss of a morpheme boundary have all occurred preferentially when the product of change is a disyllabic word. Blust suggests that conspiracies arise when a pattern in a language becomes particularly salient, leading it to function as a 'linguistic attractor' in

language change (Cooper 1999). Within variationist approaches to language change (e.g. Ohala 1989, Labov 1994, Bybee 2001, Blevins 2004 and many others), a salient pattern can influence the course of language change by biasing categorization (cf. 'Change' in Evolutionary Phonology, Blevins 2004), and/or by biasing the range of variation in production. (For a simulation of pattern feedback in production biasing language change, see Wedel (2007)). In this context, Blust notes that disyllables have been reported to make up 94% of the set of content words in proto-Austronesian, and that in many modern Austronesian languages the disyllable remains the dominant word type.

A wide range of evidence is consistent with the hypothesis that such linguistic attractors bias individual behavior and thereby influence the course of language change. Many studies have shown that variation in linguistic behavior is biased toward previous experience: both grammaticality judgements (e.g., Krott et al. 2002, Albright 2002, Pierrehumbert 2006a) and production variation/errors (e.g., Bybee & Moder 1983, Dell et al. 2000, Vitevitch & Sommers 2003, Gonnerman et al. 2007, and many more) are biased by similarity and by pattern type-frequency at a wide range of representational levels (for reviews on various of these topics, see Bybee 2001, Ernestus & Baayen 2003, Bybee & McClelland 2005, Pierrehumbert 2006b, Baayen 2007, Pisoni & Levi 2007). Patterns of simulated language change in simple model systems (Wedel 2007, 2009) are also consistent with this notion: when production and perception errors by simulated agents are biased toward previous patterns, further supporting evidence could be sought in patterns of change within iterated artificial language learning and transmission paradigms of the sort pioneered by Kirby et al. (2008). Finally, it is worth noting that within models such as Evolutionary Phonology, synchronic alternation patterns are created through diachronic change rather than through mechanisms localized within a single individual's language faculty (Blevins 2004; see chapter 3 for a review of earlier theories of this type). Under this model, this account of diachronic conspiracies provides the basis for an account of synchronic conspiracies as well.

Actuation versus propagation of change

A significant question in historical linguistics is how an initially isolated change can survive and propagate throughout a community, given that language learners tend to converge on a common community standard. To the extent that this is the case, isolated variants should never be able to gain a foothold in a speech community because every learner is exposed to many speakers (see Keller 1994, p. 99 and Nettle 1999 for discussion). In a foundational paper, Nettle (1999) uses a well-articulated simulation to explore factors that are required to allow randomly occurring variants to become established, assuming the existence of a stratified social structure. He finds that given reasonable assumptions (i.e., that *ceteris paribus*, learners tend to adopt the local majority pattern), random variation in acquisition is not sufficient to induce a population-wide transition from one pattern to another without being so pervasive as to obliterate any coherent pattern at all. He then shows that when significant prestige inequities are introduced in which a small number of individuals serve disproportionately as acquisition models, a novel variant can survive if it spreads in the population sufficiently to support itself through positive feedback. When a small number of individuals exert strongly disproportionate influence, the *effective population size* is small, allowing random events a greater chance of influencing the trajectory of change (see the literature on genetic drift in biological populations, e.g., the introductory online article by Kliman, Sheehy & Schultz (2008)). However, it is clear that functional articulatory and perceptual factors influence the course of change as well; otherwise, we should observe as many diachronic changes that are phonetically unnatural as natural. Nettle explores the influence of functional biases in his model, and concludes that in order to enforce change alone, functional biases have to be sufficiently strong that anti-functional patterns should never occur. Since this is not the case, Nettle argues that social factors are a critical engine of change, but that the rate of actuation and the efficiency of propagation must also be biased by functional factors that influence ease of production, perception and acquisition.

Emergence of phonemes and inventory structure

Vowel inventories appear to be constructed to optimize perceptual contrast between neighboring vowels given extant articulatory constraints (Liljencrants and Lindblom 1972). How does this apparent optimization come about? In an early self-organizational approach to this problem using a perception-production feedback loop, de Boer (2000) proposed that structure in vowel inventories emerges through

interaction of language users under perceptual and production constraints, assuming a tendency for language users to imitate each other (see also Browman and Goldstein 2000 for an abstractly similar self-organizational account couched within Articulatory Phonology). To test and illustrate this idea, de Boer constructed a simulation in which a group of agents can produce, perceive and remember vowel pronunciations in the form of prototypes. (Agents are entities within a simulation that can change independently, here representing individual language users.) Within the simulation, agents speak and imitate each other, modifying their vowel categories in response to how successful their imitations are. In each round, a random pair of agents is chosen from the larger set of agents to act as speaker and hearer. The speaker articulates a randomly chosen vowel from memory with some random error; if it has no vowels in memory, it produces a random vowel within the available articulatory space. The hearer compares the formant values of the vowel to prototypes it has in memory and chooses the closest one. If it has no vowels in memory, it creates a similar vowel and calculates its associated articulatory parameters. The listener then repeats the matched prototype vowel for the speaker, who checks to see how close it was to the originally produced vowel. If the vowel is judged to be the same, the speaker agent gives the listener feedback that its imitation was successful. In that case, the listener shifts the parameters of its matching prototype closer to the vowel that it heard from the speaker. If the imitation was not successful, the listener checks its memory to find out how often that prototype has given rise to successful imitations. If it has been mostly unsuccessful, it moves that prototype closer to the sound it heard, just as in a successful imitation. If it has been mostly successful before, it may be that the speaker has an additional prototype vowel in that region of vowel space, and so the listener creates a new prototype to approximately match what it heard. Several additional processes come into play in this simulation: (i) if a vowel prototype is infrequently matched to a perceived vowel, it is discarded; (ii) if two vowel prototypes are too close together, they are merged, and (iii) new vowels are introduced by speakers at a low frequency.

This simulation employs a number of features that may not correspond directly to actual features of language use (e.g., direct feedback on imitative success; the operative mechanisms of category loss and merger), but that is not its primary point. These mechanisms simply allow vowel inventories in individual agents to change over time in response to constraints on the differentiation of vowels that are perceived in individual usage events. de Boer shows that given these constraints, populations of agents starting with empty vowel inventories develop jointly held, phonetically natural vowel inventories. He concludes from this that the typological generalizations over vowel inventories found in natural language may arise through articulatory and perceptual constraints in usage rather than some more direct, innate specification. Coherent structure is primarily driven by positive feedback in this system, which comes in two forms: modification of prototypes toward perceived vowels, and merger of prototypes that get too close. These encourage the development of coherent vowel categories shared across the set of agents. Because vowels that are too perceptually confusable tend to be merged, the set of surviving vowels tends towards a perceptually 'optimal' arrangement.

Pierre-Yves Oudeyer (2002, 2006) has used an abstractly similar, more physiologically grounded model of a perception/production feedback loop to argue that positive feedback inherent in processing can create categorial distinctions in the absence of any functional pressure. Research in response biases of cortical fields of neurons shows that their output is well predicted by the aggregate response of the entire field, rather than by the output of the most highly activated neuron. From the set of activities of all neurons, it has been found that one can predict the perceived stimulus or motor output by computing the population vector over the field, namely, the sum of all preferred outputs of the set of neurons multiplied by their activities (Georgeopoulos, Schwartz, and Ketter 1986; for an account of the perceptual magnet effect (Kuhl 1991) based in this phenomenon, see Guenther and Gjaja 1996). The important feature of the population vector for our purposes is that it is shifted toward the center of the local distribution of outputs relative to the most highly activated neuron. Given a close mapping between perception and production (Oudeyer 2002, Fowler and Balantucci 2005), this property of cortical fields should produce positive feedback promoting the coalescence of perceptual-motor categories into well-defined distributions over many cycles of use.

In Oudeyer's (2002) model, linked motor and perceptual cortical fields are initially populated with randomly

tuned neurons, such that there are no distinct coherent sound categories. Over the course of the simulation, randomly chosen production stimuli are produced by the motor field and processed by the perceptual field. In processing, each neuron in the perceptual field is activated by the production stimulus under the control of a Gaussian tuning factor responsive to the degree of match between the stimulus and a neuron's preferred vector. The preferred vectors of all neurons that have been activated to some degree by the stimulus are then shifted toward that of the maximally active neuron, producing a reversion to the local mean. This update function acts to incrementally consolidate the vectors exhibited by the neural map, influenced by random peaks in the distribution of stimuli produced early in the simulation. The perceptual and motor fields are linked by an update function that shifts vectors in the motor field in parallel to those in the perceptual field, closing the perception/production feedback loop. The resulting positive feedback between perception and production allows a rapid collapse of the originally random distribution of vectors in the sensory map into a small number of coherent sound-motor categories. Oudeyer interprets this feature of his model to suggest that native features of our neurological production and perception apparatus may be designed to develop categories of a particular granularity, and that this feature may play a role in the development of phoneme inventories.

Merger versus contrast maintenance

When one sound category becomes more similar to another over the course of sound change, one possible outcome is category merger, as has occurred between /a/ and /ɔ/ in western dialects of North American English. Often however, the set of categories translates through phonetic space in a 'chain shift' such that the system of contrasts is maintained even though the specific phonetic correlates of each category change (see Gordon (2003) for a review). A number of feedback-based models have been recently proposed that provide accounts for the mechanisms of both merger and/or shift and make predictions about the conditions under which either may occur. To understand these models, it will be helpful to briefly review the role of experienced phonetic variation in production and perception. A wide variety of experimental evidence indicates that individual percepts can leave detailed, long-lived traces in memory, and that these memory 'exemplars' influence future perception and production behavior (for reviews, see Tenpenny 1995, Johnson 1997; see also Pierrehumbert 2006b and volume 23 of *The Linguistic Review*). The influence of perception on production (Goldinger 2000, Nielsen 2007) creates the possibility of a perception/production feedback loop in which the effect of biases anywhere in the cycle can potentially build up over time to shift behavior within a single generation. Janet Pierrehumbert used an exemplar-based model of this loop to explore the consequences of feedback for merger between perceptually adjacent phonological categories (2001, 2002). In this model, categories consist of an abstract label and a set of stored perceptual exemplars that have been mapped to that category, where each exemplar is associated with an activation level that decays exponentially over time. No proposed mechanism in this particular model requires transmission between distinct agents, so as a simplification the simulation architecture uses a single category system in conversation with itself. Production proceeds by probabilistically choosing an exemplar in relation to activation level, averaging all the exemplar values within a set window around the chosen exemplar in proportion to their activations, and then adding a small amount of normally-distributed noise to that average. Averaging within a window around a single exemplar creates a reversion to the mean of the local distribution, just as the use of the population vector does in Oudeyer's simulation described above. Adding noise to production outputs keeps the distribution from collapsing to a single point through the effect of averaging and allows the system to evolve over time. To decide what label the new output should be categorized under, the summed activation level of exemplars within a set window around the output value is calculated for each label. The percept is then stored as a new exemplar under the category label with the highest score.

Pierrehumbert showed that given this architecture, if two categories drift close enough such that they begin to compete for percepts along their mutual boundary, the category with greater overall exemplar activation tends to eventually absorb the less active category. This occurs through positive feedback between current activation and the ability to compete for percepts. All else being equal, an ambiguous percept is more likely to be mapped to a more active category than a less active category, which only results in the active category becoming yet more active with respect to the less active category. This snowballing feedback results in

more and more percepts being mapped to the more active category, until the activation of the other category eventually falls low enough that it effectively no longer exists and is absorbed into the more frequent category. An example that can be modeled this way is the above mentioned merger of /ɔ/ into the more frequent /ɑ/ in western dialects of American English.

In all of the work reviewed above, local, similarity-based positive feedback drives coalescence of system elements into categorial groupings. In none of these accounts, however, is there any repulsive force that would prevent the steady merger of categories over time as they eventually drift into one another. As a consequence, the maintenance of multiple categories over the course of language change in these models would require either regular generation of new distinctions as in the de Boer (2000) model above or some mechanism to favor preservation of at least some existing contrasts. Paul Boersma and Silke Hamann (2008) have approached the problem of contrast maintenance between existing sound categories through a constraint-based model that makes use of categorization accuracy on the part of a language learner. As a demonstration of their model, Boersma and Hamann simulate the evolution of category label/contents mappings within a unidimensional space. To concretize the model, they use the spectral frequency range of sibilants in human languages as the perceptual space. (In the following brief discussion of their model, I use an /s/ + /ʃ/ two-sibilant system as a running example, although more or fewer categories are possible.) The architecture of the model is vertical, in which a naive agent learns to associate part of the spectral frequency range with /s/ and another part with /ʃ/ by hearing examples from a teaching agent with feedback. After this learning phase, the agent becomes a teacher and produces examples of /s/ and /ʃ/ to a new learner agent, and so on. For the purposes of the argument, Boersma and Hamann assume that learning agents have acquired sound category labels from word patterns prior to the beginning of the simulation. As a result, learners know at the start how many sibilant categories their language has, but are not yet sure where their distributions lie within the frequency continuum. In addition, learners have the ability to learn an association between a given spectral frequency and /s/ or to /ʃ/ by constructing a ranking among Optimality Theoretic constraints banning the mapping of a particular frequency to a particular category label (Boersma 1997). The architecture of the perception grammar incorporates both frequency of presentation and categorization accuracy, with the result that the grammar is maximally 'certain' about mappings for sibilant frequencies that are further apart relative to those frequencies that were most often heard. Because subsequent production is based on sampling from the learned perception grammar rather than from the distribution of actually learned examples, an agent's production favors a distribution of sibilants that is slightly better separated than that which she heard herself. This creates a positive feedback loop that promotes increasing contrast between categories over many teacher/learner cycles. However, agents' productions are also influenced by ranked articulatory constraints that have the effect of biasing productions toward the center of the frequency continuum. Boersma and Hamann show that under the balancing influence of positive feedback via the perception grammar and negative feedback from articulatory constraints, categories evolve as well-spaced distributions with a joint center of gravity at the midpoint of the continuum.

The Boersma and Hamann model relies on error-feedback from the lexicon within the learner to drive contrast-maintenance. Another model for contrast-maintenance has been proposed that operates in situations in which there is an *absence* of error-feedback (Wedel 2004, Blevins and Wedel 2009; see also Ettliger 2007). The relevant absence of error feedback in this model occurs when an ambiguous pronunciation is not rescued by an external context to determine its intended category mapping. As an example of external disambiguation, words like 'too' and 'two' in English are rarely confused by listeners because they are used in distinct sentential and semantic contexts. In contrast, because of their semantic similarity, words within morphological paradigms are often distinguished in context primarily by their phonetic differences. For example, utterances like 'I cook chicken well' and 'I cooked chicken well' could be used in very similar contexts, in which case the tense of the verb 'cook' is conveyed almost entirely by the audible presence or absence of the past-tense [-t].

A hypothesis introduced in Wedel (2004) and explored more deeply in Blevins and Wedel (2009) is that this effectively greater 'functional load' of word-internal phonetic information within paradigms may account for anti-homophony effects in paradigms. 'Anti-homophony' refers to the failure of otherwise regular sound changes to occur in words when that change would render them homophonous with another word (see

Blevins and Wedel 2009 for a review and examples). As in the Pierrehumbert (2001, 2002) model of category-merger reviewed above, this model rests on evidence that category behavior is updated by experience: if a pronunciation is ambiguous in context, a hearer may map it to a category that the speaker did not intend, resulting in the effective trading of a variant between the two categories at their boundary. It is this ‘variant trading’ between perceptually adjacent categories that drives the behavior of the model by preserving a crisp boundary between adjacent categories (see Blevins and Wedel 2009 for more discussion).

Looking forward

Our ability to build theories is limited by the knowledge that we already have. Whether or not any of the current self-organization-based accounts in phonology are ‘right’, they are valuable in expanding our understanding of pattern-formation mechanisms. Self-organization is ubiquitous in physical, biological and cultural systems, and given that language provides the conditions for self-organization many times over, linguists should anticipate finding it as a contributing mechanism in this domain as well. Just as with any other type of account, however, showing that a particular structure could arise through self-organization does not mean that it does so. A hypothesis stands or falls on the empirical success of its predictions. Arguments from first principles about, for example, whether phonological patterns are more likely to derive from a language-specific cognitive faculty or a more general set of factors are less valuable than well-constructed tests of model predictions. Fortunately, a wide variety of techniques and approaches are now available for testing hypotheses, from the ever-growing array of psycholinguistic techniques through corpus studies and artificial language learning paradigms.

Self-organizational models for structure formation make use of previously identified cognitive, articulatory, perceptual or social factors as contributing building blocks. In turn, these models make further predictions about those factors, or may predict some yet-undescribed phenomenon. For example, self-organizing models of phonological change through usage require a production/perception feedback loop that can drive small, but persistent and generalizable changes in post-acquisition phonological categories (e.g., Bybee 2001, Pierrehumbert 2001, 2002, Wedel 2007, Mielke 2008). Although more work needs to be done to establish their generality, results from a variety of psycholinguistic studies are consistent with this prediction (e.g., Goldinger 2000, Nielsen 2007, Kraljic and Samuel 2006). Likewise, each of the other models reviewed above makes new predictions that can be tested empirically. As phonologists bootstrap back and forth between model building and simulation on the one hand and empirical methods on the other, the field should gain a steadily better sense of if and how self-organizational mechanisms contribute to the wide variety of phenomena that we study.

References

- Albright, Adam. 2002. Islands of reliability for regular morphology: Evidence from Italian. *Language* 78: 684-709.
- Ball, Phillip. 1999. *The Self-Made Tapestry: Pattern Formation in Nature*. Oxford University Press.
- Baayen, R. Harald. 2007. Storage and computation in the mental lexicon. In Gonia Jarema and Gary Libben, eds., *The Mental Lexicon: Core Perspectives*. Elsevier, 81-104.
- Beckman, Mary, and Jan Edwards 2000. The Ontogeny of Phonological Categories and the Primacy of Lexical Learning in Linguistic Development. *Child Development* 71: 240-249.
- Beckner, Clay, Richard Blythe, Joan Bybee, Morten H. Christiansen, William Croft, Nick C. Ellis, John Holland, Jinyun Ke, Diane Larsen-Freeman, and Tom Schoenemann. 2009. Language is a complex adaptive system. *Language Learning* 59:4S1.
- Blevins, Juliette. 2004. *Evolutionary Phonology: The emergence of sound patterns*. Cambridge: Cambridge University Press.

- Blevins, Juliette, and Wedel, Andrew. 2009. Inhibited sound change: An evolutionary approach to lexical competition. *Diachronica*, 26: 143-183.
- Blust, Robert. 2007. Dyllabic attractors and anti-antigemination in Austronesian sound change. *Phonology* 24 : 1-36.
- Bod, Rens, Jennifer Hay, and Stefanie Jannedy eds 2003. *Probabilistic linguistics*. Cambridge, MA: The MIT Press.
- Boersma, Paul and Hamann, Silke. 2008. The evolution of auditory dispersion in bidirectional constraint grammars. *Phonology* 25, 217-270.
- Browman, Catherine & Louis Goldstein. 2000. Competing constraints on intergestural coordination and self-organization of phonological structures. *Bulletin de la Communication Parlée*, 5:25-34.
- Bybee, Joan. 2001. *Phonology and language use*. Cambridge: Cambridge University Press.
- Bybee, Joan. 2002. Word frequency and context of use in the lexical diffusion of phonetically conditioned sound change. *Language Variation and Change* 14: 261-290.
- Bybee, Joan, and Carol Lynn Moder. 1983. Morphological classes as natural categories. *Language* 59. 251-270.
- Bybee, Joan, and McClelland, James L. 2005. Alternatives to the Combinatorial paradigm of linguistic theory based on domain general principles of human cognition. *The Linguistic Review* 22: 381-410.
- Camazine, Scott, Jean-Louis Deneubourg, Nigel Franks, James Sneyd, Guy Theraulaz, and Eric Bonabeau (2001). *Self-Organization in Biological Systems*. Princeton: Princeton University Press.
- Cooper, David. L. 1999. *Linguistic Attractors: The cognitive dynamics of language acquisition and change*. Amsterdam/Philadelphia: John Benjamins Publishing.
- Croft, William. 2000. *Explaining language change: An evolutionary approach*. Harlow, Essex: Longman.
- Cziko, Gary. 1995. *Without Miracles: Universal Selection Theory and the Second Darwinian Revolution*. Cambridge, MA: MIT Press.
- Dawkins, Richard 1986. *The Blind Watchmaker*. New York: W. W. Norton and Company.
- Deacon, Terence W. 1997. *The Symbolic Species: The Co-evolution of Language and the Brain*. W.W. Norton.
- de Boer, Bart. 2000 Self-organization in vowel systems. *Journal of Phonetics* 284:441-465.
- Dell, Gary S., Kristopher D. Reed, David R. Adams, and Antje S. Meyer. 2000. Speech errors, phonotactic constraints, and implicit learning: A study of the role of experience in language production. *Journal of Experimental Psychology: Learning, Memory, and Cognition* 26: 1355-1367.
- Dennet, Daniel. 1995. *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. New York: Simon and Schuster.
- Elman, Jeff L. 1995. Language as a dynamical system. In Robert Port and Timothy van Gelder, eds., *Mind as Motion: Dynamical Perspectives on Behavior and Cognition*. Cambridge, MA: MIT Press.
- Ernestus, Mirjam & R. Harald Baayen. Predicting the unpredictable: Interpreting neutralized segments in Dutch. *Language*, 79(1):5-38, 2003.
- Ettlinger, Marc. 2007. Shifting categories: An exemplar-based computational model of chain shifts.

Proceedings of the 29th Annual Conference of the Cognitive Science Society. Austin, Texas: Cognitive Science Society. 239-244.

Ferguson, Charles A. and Farwell, Carol B. 1975. Words and sounds in early language acquisition. *Language*, 51: 419–439. Reprinted with Appendix in William S.-Y. Wang, ed., 1977. *The Lexicon in Phonological Change*. The Hague: Mouton. 7–68.

Fikkert, Paula and Clara C. Levelt. To appear. How does place fall into place? The lexicon and emergent constraints in the developing phonological grammar. In Peter Avery, B. Elan Dresher, and Keren Rice, eds., *Contrast in phonology: Perception and Acquisition*. Berlin: Mouton.

Fikkert, Paula. 2007. Acquiring phonology. In Paul de Lacy, ed., *Handbook of Phonological Theory*. Cambridge, MA: Cambridge University Press. 537–554.

Fowler, Carol A. & Bruno Balantucci. 2005. The relation of speech perception and speech production. In David B. Pisoni & Robert E. Remes, eds. *The handbook of speech perception*, pp. 633-52. Malden, MA: Blackwell.

Gardner, Martin. 1970. Mathematical Games. *Scientific American*. 223: 120-123.

Garrett, Anrew. 2008. Paradigmatic Uniformity and Markedness. In Jeff Good, ed. *Linguistic Universals and Language Change*. Oxford University Press. 125-143.

Gell-Mann, Murray. 1992. Complexity and complex adaptive systems. In John A. Hawkins and Murray Gell-Mann, eds., *The Evolution of Human Languages*. New York: Addison-Wesley. 3–18.

Georgeopoulos, Apostolos, Schwartz, Andrew B., and Ketter, Ronald E. 1986. Neuronal population coding of movement direction. *Science* 233,1416–1419.

Goldinger, Stephen. D. 2000. The role of perceptual episodes in lexical processing. In Anne Cutler, James M. McQueen, and R. Zondervan, eds., *Proceedings of SWAP Spoken Word Access Processes*, Nijmegen, Max-Planck-Institute for Psycholinguistics.

Gonnerman, Laura M., Seidenberg, Mark S., and Andersen, Elaine S. 2007. Graded semantic and phonological similarity effects in priming: Evidence for a distributed connectionist approach to morphology. *Journal of Experimental Psychology: General* 136, 323-345.

Gordon, Matt. J. 2002 Investigating Chain Shifts and Mergers. In J.K Chambers, Peter Trudgill and Natalie Schilling-Estes, eds., *Handbook of Language Variation and Change*. Blackwell Publishers Ltd, MA, 244-266.

Guenther, Frank H., and Gjaja, Marin. N. 1996. The perceptual magnet effect as an emergent property of neural map formation. *Journal of the Acoustical Society of America* 100, 1111-1121.

Hare, Mary, and Elman, Jeff L. 1995. Learning and morphological change. *Cognition* 56, 61-98.

Haspelmath, Martin. 1999. Optimality and diachronic adaptation. *Zeitschrift für Sprachwissenschaft* 18:180–206.

Hock, Hans H. 2003. Analogical Change. In: Brian Josephs and Richard Janda, eds., *The Handbook of Historical Linguistics*. Blackwell 441-480.

Hurford, James. 1999. The evolution of language and languages. In Robin Dunbar, Chris Knight, and Camilla Power, eds., *The Evolution of Culture*. Edinburgh: Edinburgh University Press.

Johnson, Keith. 1997. Speech perception without speaker normalization. In Johnson, Keith. and Mullenix John W., eds., *Talker Variability in Speech Processing*. San Diego: Academic Press.

Kauffman, Stuart. 1995. *At Home in the Universe: The Search for the Laws of Self-Organization and*

Complexity. Oxford University Press.

Keller, Rudi. 1994 *On Language Change: The Invisible Hand in Language* London: Routledge.

Kirby, Simon. 1999. *Function, Selection and Innateness: The Emergence of Language Universals*. Oxford: Oxford University Press.

Kirby, Simon, Hannah Cornish & Kenny Smith. 2008. Cumulative Cultural Evolution in the Laboratory: an experimental approach to the origins of structure in human language. *Proceedings of the National Academy of Sciences* 10531:10681-10686.

Kisseberth, Charles. 1970. On the functional unity of phonological rules. *Linguistic Inquiry* 1:291-306.

Kliman, Richard, Bob Sheehy, and Joanna Schultz. 2008. Genetic Drift and Effective Population Size. *Nature Education* 13. <http://www.nature.com/scitable/topicpage/Genetic-Drift-and-Effective-Population-Size-772523>

Kraljic, Tanya, and Arthur Samuel. 2006. Generalization in perceptual learning for speech. *Psychonomic Bulletin and Review* 13, 262-268.

Krott, Andrea, Robert Schreuder, & R. Harald Baayen. 2002: Analogical hierarchy: Exemplar-based modeling of linkers in Dutch noun-noun compounds. In Royal Skousen, Deryle Lonsdale, and Dilworth B. Parkinson, eds., *Analogical Modeling: An Exemplar-Based approach to Language*. Amsterdam: John Benjamins. 181-206.

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.

Liljencrants, Johann, and Björn Lindblom. 1972: Numerical simulation of vowel quality systems: The role of perceptual contrast. *Language* 48: 839-862.

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

Lindblom, Björn. 1992. Phonological units as adaptive emergents of lexical development. In: Carol A. Ferguson, Lise Menn, and Carol Stoel-Gammon, eds., *Phonological Development: Models, Research, Implications*. Timonium: York Press. 131-163.

Lindblom, Björn, Peter MacNeilage & Michael Studdert-Kennedy. 1984. Self-organizing processes and the explanation of language universals. In Brian Butterworth, Bernard Comrie, and Östen Dahl, eds., *Explanations for Language Universals*. Walter de Gruyter & Co.

MacWhinney, Brian. 2006 The emergence of linguistic form in time. *Connection Science* 17, 191-211.

Medeiros, David. 2008. Optimal Growth in Phrase Structure. *Biolinguistics* 2: 156-195.

Mielke, Jeff. 2008. *The emergence of distinctive features*. Oxford: Oxford University Press.

Nettle, Daniel. 1999 Using Social Impact Theory to simulate language change. *Lingua* 1082-3:95--117.

Nielsen, Kuniko. 2007. The Interaction between Spontaneous Imitation and Linguistic Knowledge. *UCLA Working Papers in Phonetics* 105, 125-137.

Lindblom, Björn. 2000. Developmental origins of adult phonology: the interplay between phonetic emergents and the evolutionary adaptations of sound patterns. *Phonetica*, 57, 297-314.

Ohala, John. 1989. Sound change is drawn from a pool of synchronic variation. In Leiv E. Breivik and

- Ernst, H. Jahr, eds., *Language Change: Contributions to the study of its cause*. Berlin: Mouton de Gruyter.
- Oudeyer, Pierre-Yves. 2002. Phonemic coding might be a result of sensory-motor coupling dynamics. In Bridget Hallam, Dario Floreano, John Hallam, Gillian Hayes, and Jean-Arcady Meyer, eds., *Proceedings of the 7th International Conference on the Simulation of Adaptive Behavior*. Cambridge: MIT Press. pp. 406-416.
- Oudeyer, Pierre-Yves. 2006 *Self-Organization in the Evolution of Speech*. Oxford University Press.
- Peck, Steven L. 2004. Simulation as experiment: a philosophical reassessment for biological modeling. *Trends in Ecology and Evolution* 19: 530-534.
- Phillips, Betty S. 1984. Word frequency and the actuation of sound change. *Language* 60:320-342.
- Phillips, Betty S. 2006. *Word frequency and lexical diffusion*. Basingstoke: Palgrave Macmillan.
- Pierrehumbert, Janet. 2001. Exemplar dynamics: Word frequency, lenition, and contrast. In Joan Bybee and Paul Hopper, eds., *Frequency Effects and the Emergence of Linguistic Structure*. John Benjamins, Amsterdam. 137-157.
- Pierrehumbert, Janet. 2001. Exemplar dynamics: Word frequency, lenition, and contrast. In Joan Bybee and Paul Hopper, eds., *Frequency Effects and the Emergence of Linguistic Structur*. John Benjamins, Amsterdam, 137-157.
- Pierrehumbert, Janet. 2002. Word-specific phonetics. In Carlos Gussenhoven and Natasha Warner, eds., *Laboratory Phonology 7*. Berlin; New York : Mouton de Gruyter.
- Pierrehumbert, Janet. 2003. Phonetic Diversity, Statistical Learning, and Acquisition of Phonology. *Language and Speech* 46: 115-154.
- Pierrehumbert, Janet. 2006a The Statistical Basis of an Unnatural Alternation, in Louis Goldstein, Doug H. Whalen, and Catherine Best, eds., *Laboratory Phonology VIII, Varieties of Phonological Competence*. Mouton de Gruyter, Berlin, 81-107.
- Pierrehumbert, Janet. 2006b. The new toolkit. *Journal of Phonetics* 34: 516-530.
- Pisoni, David. B. and Levi, Susannah V. 2007. Some observations on representations and representational specificity in speech perception and spoken word recognition. In Gareth Gaskell, ed., *The Oxford Handbook of Psycholinguistics*. Oxford University Press. 3-18.
- Plaut, David C. and Christopher T. Kello. 1999. The Emergence of Phonology From the Interplay of Speech Comprehension and Production: A Distributed Connectionist Approach. In Brian MacWhinney, ed., *Emergence of Language*. Hillsdale, NJ: Lawrence Earlbaum Associates.
- Prince, Alan and Paul Smolensky 1993. Optimality Theory: Constraint interaction in generative grammar. Technical Report CU-CS-696-93, Department of Cognitive Science, University of Colorado at Boulder, and Technical Report TR-2, Rutgers Center for Cognitive Science, Rutgers University, New Brunswick, NJ.
- Steels, Luc. 2000. Language as a Complex Adaptive System. In Marc Shoenauer, ed., *Proceedings of PPSN VI, Lecture Notes in Computer Science*. Berlin: Springer-Verlag. 17-26.
- Stemberger, Joseph, and MacWhinney, Brian. 1988. Are inflected forms stored in the lexicon? In Michael Hammond and Michael Noonan, eds., *Theoretical Morphology*. New York: Academic Press.
- Tenpenny, Patricia L. 1995. Abstractionist versus episodic theories of repetition priming and word identification. *Psychonomic Bulletin and Review* 2: 339-363.

- Vihman, Marilyn M., and Croft, William. 2007. Phonological development: Toward a 'radical' templatic phonology. *Linguistics*, 45: 683–725.
- Vihman, Marilyn, DePaolis, Rory. A., Keren-Portnoy, Tamar. 2009. A dynamic systems approach to babbling and words. In Edith. L. Bavin, ed., *The Cambridge Handbook of Child Language*. Cambridge: Cambridge University Press. 163-182.
- Vitevitch, Michael S. and Sommers, Mitchell. 2003. The facilitative influence of phonological similarity and neighborhood frequency in speech production. *Memory and Cognition* 31:491-504.
- Wang, William S.-Y. 1969. Competing changes as a cause of residue. *Language* 45:9–25.
- Wang, William S.-Y., and Cheng, Chin-Chuan. 1977. Implementation of phonological change: The Shaungfeng Chinese case. In William. Wang, ed., *The Lexicon in Phonological Change*. The Hague: Mouton. 86–100.
- Wedel, Andrew. 2004. Category competition drives contrast maintenance within an exemplar-based production/perception loop. In John Goldsmith, and Richard Wicentowski, eds., *Proceedings of the Seventh Meeting of the ACL Special Interest Group in Computational Phonology* 7: 1-10.
- Wedel, Andrew. 2007. Feedback and regularity in the lexicon. *Phonology* 24: 147-185.
- Wedel, Andrew 2009. Variation, multi-level selection and conflicts between phonological and morphological regularities. In Juliette Blevin and James Blevins, eds., *Analogy in Grammar: Form and Acquisition*. Oxford University Press.