

Chaotic Evolution in an Unbiased Learner

Adam Albright

Massachusetts Institute of Technology

1. Introduction

A great deal of research in theoretical phonology has been motivated by a single fundamental question: what kinds of grammatical restrictions or biases are needed in order to explain the observed phonological typology? At least three major positions can be discerned in the recent literature. One possibility, assumed in much current work in Optimality Theory, is that grammar consists of a fixed innate constraint set allowing more or less exactly the set of attested languages (Prince and Smolensky, 1993/2004). Under this approach, unattested languages are completely unlearnable. A more moderate position is that grammar may capture a wide variety of natural and unnatural patterns, but it is biased towards certain “natural” patterns—perhaps by favoring phonetically grounded constraints (Hayes, 1999; Hayes and Steriade, 2004; Wilson, 2006). Under such an approach, unattested patterns may be learnable, but correspond to dispreferred grammars. Finally, some authors hold that the grammatical mechanism is not constrained at all by phonetic considerations, but that certain structures are unattested for “extra-grammatical” reasons. For example, an unattested structure may be difficult to produce or perceive accurately and is thus lost across generations (Ohala, 1981 et seq; Blevins, 2004; and others), or the pattern may be rare because it requires an unlikely conspiracy of multiple independent changes to establish it (Harris, 2008).

Given that typology may be shaped by both grammatical and extra-grammatical factors, there has been a recent trend to move away from explaining typological asymmetries, seeking instead “real-time” evidence that humans show biases towards certain phonological patterns over others (see, e.g., Pycha, Nowak, Shin, and Shosted, 2003; Wilson, 2003; Seidl and Buckley, 2003; Kawahara, 2006; Peperkamp, Calvez, Nadald, and Dupoux, 2006; Wilson, 2006; Zhang and Lai, 2006; Becker, Ketrez, and Nevins, 2007; Berent, Steriade, Lennertz, and Vaknin, 2007; Moreton, 2007; Zuraw, 2007; Finley and Badecker, in press). An equally important line of inquiry, however, is to pursue formal models that combine grammatical and extra-grammatical factors to derive typological asymmetries. This paper represents a small step in this direction, employing computational simulations to investigate the types of grammatical constraints or biases that may be needed for a typological tendency to emerge.¹

Specifically, we seek to assess the claim that typologically common patterns emerge naturally as the diachronic end state of a series of phonetically natural misproductions and/or misperceptions (Ohala, 1981; Blevins, 2004; Blevins & Garrett, 2004). This claim can be represented schematically as follows: suppose that in an initial state, we start with a language that has a typologically disfavored nasal place contrast before stops: [ampa] vs. [anpa]. In such a language, we would expect coarticulatory overlap to cause some realizations of /anpa/ to be realized with a certain degree of labialization: [an̠mpa], [aⁿmpa]. Substantially coarticulated [an̠mpa] tokens may be difficult to distinguish from [ampa] tokens (as the intended realization of /ampa/) (Hura, Lindblom, and Diehl, 1992; Steriade, 2001; but see also Winters, 2001). If at least some of these coarticulated tokens are categorized by listeners as [ampa] rather than as intended [anpa], then over time the number of /np/ words may dwindle, leaving subsequent generations with less and less support for the contrast. A language with such an asymmetry is shown in Table 1. Furthermore, once listeners come to expect nasals to be labial before labial stops, this statistical bias may lead to a perceptual bias towards agreeing nasals, further reinforcing the pattern (Bybee, 2001; Pierrehumbert, 2001). The end result, then, would be a language with the typologically common pattern of nasal place assimilation. Crucially, however, in this account the grammar has no built-in preference for nasal place assimilation (*np). If a language were to somehow instead lose all coronal nasals before

Thanks to Edward Flemming, Bruce Hayes, Michael Kenstowicz, Donca Steriade, Colin Wilson, and the audience at WCCFL 27 for helpful comments and discussion; all remaining errors are, of course, my own.

¹See Morley (2008) for a closely related line of inquiry.

Table 1: Preference for nasal place agreement, set up by misperception of intended [np] as [mp]

	/n/	/m/
___ p	ronpa, tunpa	ampil, sampo, rimpa, banmplo, nampa, limpio, gampita, bampet, gempas, krampuk, wampa
Elsewhere	tona, rino, kano, noli, stana, minu	mafa, maku, poma, rami, morta, masu

fricatives (*ns), the learner would have just as easily been able to learn a pattern of labialization by fricatives (/ns/ → [ms]).

The idea that statistical asymmetries may be established and reinforced through perception has intuitive appeal, and the power of this mechanism has been computationally demonstrated for simple cases (Kochetov, 2004; Wedel, 2007). These simulations have been quite limited in the types of patterns that they can encode, however. Typically, they have involved a contrast with no conditioning context, or conditioned by just a single neighboring segment. The dynamics of reinforcement in a statistical learning model may be quite different if the model is able to encode more complex patterns. To see why, consider again the language in Table 1. As noted above, this language shows a strong statistical tendency for nasals to agree in place with a following stop. However, on closer inspection, a more subtle generalization is also possible: nasals followed by [p] are coronal if preceded by a round vowel (*ronpa*) and labial otherwise (*ampil*). If this vowel effect is noticed, it could guide expectations and be reinforced through subsequent perceptual categorization, leading to a phonetically unnatural allophonic distribution. In point of fact, I know of no language in which the outcome of nasal place “assimilation” is sensitive to the preceding vowel.

How could such accidental correlations arise? Those who assign phonology problem sets frequently encounter this phenomenon: irrelevant correlations are abundant in limited data sets, and assiduous learners with no particular expectations eagerly latch on to them. Such correlations are highly unlikely in larger data sets, but unfortunately language data is often sparse, and due to sampling limitations the learner may not have access to large numbers of words containing round vowel + nasal + [p] sequences. Furthermore, gradient statistical patterns due to change in progress may exacerbate this problem by making data concerning a particular sequence even sparser, increasing the chances that the remaining [mp] examples all happen to share some other property in common as well. This constitutes a type of poverty of the stimulus issue: although the underlying mechanism of change was a simple one, the limited set of examples available to the learner may side-track learners to arrive at an analysis that is complex and unlike any attested pattern. I will refer to such complications as “chaotic” evolution.

The goal of this paper is to assess how serious a problem chaotic evolution is for diachronic accounts of typological asymmetries. As the language in Table 1 shows, it is easy to construct schematic examples in languages with just a few words, but this gives little indication of how likely such situations would be in more realistic cases with larger vocabularies. The strategy is to explore the predictions of a relatively powerful inductive learner that incorporates minimal substantive grammatical biases, to see what kinds of patterns it can discover from randomly generated lexicons in which perceptually difficult sequences are somewhat rare. As we will see in section 2, accidental correlations appear to arise relatively frequently, raising the danger that phonetically motivated sound change might often be diverted to create unnatural phonological distributions. In section 3, I discuss some grammatical biases that may help to avoid such situations. Importantly, although phonetically guided markedness biases such as nasal place assimilation are one possible solution, it turns out that other, simpler grammatical constraints are equally effective in the current case. Thus, although this study motivates the need for some sort of constraints or biases on inductive learning, identifying the exact nature of such constraints is left as a matter for future research.

2. Unbiased learning of transitional languages

As a concrete test case for unbiased learning, we consider the extremely common pattern of nasal place assimilation: nasals (and especially /n/) must be homorganic with following obstruents. As

described above, it is very plausible that this pattern could come about through coarticulation and reanalysis, under the assumption that coarticulation of [n] with a following labial stop could occasionally cause (intended) [np] tokens to be perceived as [mp], gradually introducing the appearance of a process of place assimilation, and perhaps also reducing the number of /np/ words. We attempt to assess here the probability that other coincidental subregularities could develop—e.g., assimilation blocked by the surrounding vowels, nearby laterals, the presence of a following coda, and so on.

The starting point for testing the probability of such subregularities was a set of artificial /VNpV/ languages. Each of these languages contained both /np/ and /mp/ clusters, embedded in words of one to four syllables. Syllable onsets drawn with equal probability from a set containing both singletons and clusters (yielding approximately 24% probability of onset clusters): {p, t, k, b, d, g, f, s, v, z, m, n, l, r, w, j, pr, tr, kr, pl, kl}. Nuclei were drawn from the set {i, e, a, o, u}. Codas were optional, occurring in approximately 35% of syllables, and were drawn from the same set of consonants as the singleton onsets. All words had at least one syllable, and the probability of each additional syllable was 33%. An example lexicon constructed in this fashion is shown in (1).

(1) Sample lexicon for VNpV language

vat, pan, kropri, ro, jeppro, konu, mij, trup, tre, bi, vezmolu, jo, se, pla, me, pa, ze, loru, japidu, buba, ma, zan, tra, lan, raple, vipra, la, do, ta, pri, klojtufe, wine, paf, trizfu, bitra, nola, ba, zada, po, ga, ...

Since the current interest is in learning the distribution of [n] and [m] before [p], we focus here on the subset of each lexicon containing [mp] and [np] sequences. One thousand VNpV languages were constructed with 50 /Np/ words each. In order to simulate languages partway through a change of [np] > [mp], these languages were constrained to have [mp] words outnumber [np] words, in a ratio of approximately 3:1. For lexicons of 50 /VNpV/ words, this translated to 37–38 /mp/ words and 12–13 /np/ words—i.e., a strong but not absolute tendency for coronal nasals to assimilate to a following labial stop. A sample of the relevant portion of such a lexicon is shown in (2). Note that since we are focusing here on words with medial /Np/ clusters, the words in this portion of the vocabulary look somewhat “cluster-heavy”.

(2) Sample lexicon: 50 /Np/ words

np	mp		
nenpa	pempi	momplo	maplompes
bonpi	jompa	dimpu	kitimpi
punpi	kanimpru	gimprapte	lepempli
kranpa	gempi	sampal	limpoflu
ganplo	primplaw	femplo	kempakro
kranpu	simpe	namprospe	fampuvwi
denpro	pimpebo	prompa	plugramplud
minpiw	wumpe	jimpav	plumpegi
kinpog	dompo	samplapem	gampagfal
mufonpa	zamplef	wemplupkro	trompepru
junpezo	zumpa	lompig	tumputro
tenpelep	pompajik	sompup	
kleblenpe	lampa	wompafa	

In all 1,000 of these languages, there is a strong tendency for nasal place assimilation. The question of interest, though, is how many of these languages contain incidental regularities governing the distribution of [n] vs. [m] in Np clusters. In order to test this, all 1,000 languages were fed to a “minimally biased” inductive model (Albright & Hayes, 2006), containing a general apparatus for constructing phonological grammars (constraints, stated in terms of phonological representations with features), but no built-in biases for grammars reflecting articulatory or perceptual ease, naturalness of interactions between targets and triggers, simplicity of contexts, generality, economy, and so on. The task of the model was to learn a grammar that could predict, to the extent possible, the distribution of

[n] and [m] in each language. Intuitively, since the following [p] is responsible for the misperceptions that have created the statistical imbalance in the first place, we expect that it should be the best predictor of nasal place—i.e., a 75% chance of [m] before [p]. However, if better predictors happen to exist, the learner may discover them as well, as particularly reliable phonological subregularities (“islands of reliability”: Albright, 2002; Albright & Hayes, 2003).

For each language, the training data fed to the model was the set of words, marked for the location of the nasal within the word and its place of articulation. Since the model was developed to learn the relation between pairs of morphologically related forms (see Albright & Hayes, 2006, for details), pairs were constructed in which the “input” contained an indication of the location of the nasal, and the “output” contained the place of articulation: (boNpi → bonpi, peNpi → pempi). Thus, the task of the model was to construct a grammar that takes underlying forms with nasals unspecified for place, and attempts to predict the place of articulation of the nasal given the surrounding context. In order to do this, the model compares pairs of words with the same nasal place (N→n or N→m), to discover the most reliable conditioning contexts for [n] vs. [m]. Individual pairs of words are compared by finding the material that they have in common, as illustrated in (3).

(3) Comparison procedure for extracting shared material

	b	o	N→n	p	i	
+	g	a	N→n	p	l	o
=	[-son -cont +voi]	[+syl +back -high]	N→n	p	(l?)	[+syl -low]
+	n	e	N→n	p		a
=	[-cont +voi]	[+syl -high]	N→n	p	(l?)	[+syl]

The comparison procedure, in brief, is as follows: first, the optimal phonetic alignment between the segments of the two words is calculated (i.e., the minimum string edit distance; Kruskal 1983), with the phonetic similarity of two segments defined according to the number of shared natural classes (Frisch et al., 2004). This alignment is represented visually in (3) as vertical alignment of segments into columns. Once the optimal alignment between two words has been established, corresponding segments are compared to extract the features that matched segments have in common, shown here in bracket notation. When one word contains material with no correspondent in the other word, it is marked as optional, shown here with parentheses and question marks. When such comparisons are iterated over a diverse set of words, the model discovers ever more general contexts. In the present case, given a large enough set of randomly generated words, we expect the learner to discover that words with [n] or [m] share nothing in common other than the following [p].² Given limited data, however, most or all examples may coincidentally happen to share other features as well. The question here is how often this occurs, among lexicons that were not designed to contain any such patterns.

In order to test the prevalence of emergent generalizations, the learning model was trained on all 1,000 artificial languages. Recall that for all of these languages, *np is satisfied 75% of the time, meaning that overall, there is only a 25% chance of the nasal ‘N’ being realized as [n]. Seen from this point of view, any context that the learner can find which correctly characterizes more than 25% of the existing [np] words is an island of reliability for [np], which could potentially help to stabilize or preserve [n] in this context. However, since the logic of feedback and reinforcement is that listeners may become biased to preferentially perceive certain segments in certain contexts, I take here a more conservative estimate of what constitutes a strong predictor of [n]—namely, any context favoring [np] with greater than 50% reliability.

The results of the simulations show that in 678 out of the 1,000 artificial languages, a context could be found that favored [np] over [mp] with better than 50% reliability. These contexts are often quite complex. For example, in language #715, the constraint in (4) correctly characterizes every single [np] example.

²As well as the inviolable ... C(liquid)VnpV... syllable structure shared by all words in the constructed artificial languages.

$$(4) \quad *m/\# \left[\begin{array}{c} -\text{nasal} \\ -\text{lateral} \end{array} \right]_0 \text{ — } p \left[\begin{array}{c} -\text{nasal} \\ -\text{dorsal} \end{array} \right]_0 \#$$

“Nasal place must be coronal if preceded exclusively by non-nasal segments (other than *l*), and followed only by non-nasal segments (other than dorsals)”

In most other cases, the emergent generalization covered more than half but less than 100% of the relevant cases: 71 of the languages had a context that favored [n] in more than 75% of the [np] words, 338 languages had a context that covered more than 58% of [np] words, and 678 had a context favoring [n] in more than 50% of [np] words. Thus, unlike the schematic example in Table 1, most of these languages had generalizations that covered the remaining [np] words with less than 100% reliability. The fact that emergent generalizations are not perfect does not shield us from their effects, however, since there is abundant evidence that speakers can notice and extend statistical trends with less than 100% coverage (Zuraw, 2000; Albright & Hayes, 2003; Ernestus & Baayen, 2003; Pierrehumbert, 2006).

What are the consequence of these islands of reliability for [np]? First, they may be extended to novel strings, created through morpheme concatenation or entering the language through loanword adaptation. For example, the constraint in (4) could lead to alternations like those in (5), in which nasals assimilate in place to a following affix only if there is a preceding lateral or nasal.

(5) Contextually restricted alternations

/ulaN+pa/	→ [ulampa]	(<i>m</i> licensed by preceding lateral)
/nutiN+pa/	→ [nutimpa]	(<i>m</i> licensed by preceding nasal)
/siN+pa/	→ [sinpa]	(N→ <i>n</i> otherwise)

More important for present purposes, these highly reliable contexts for [np] could potentially help to preserve [np] by reinforcing perceptual biases. It is well known that context strongly influences categorization and identification of acoustically ambiguous tokens. Such biases may be based on the entire rest of the word: for example, ambiguous tokens are preferentially identified as phones that yield existing words, so that ambiguous [k]~[g] is more likely to be identified as /g/ in the context /___ift, since *gift* is a word while **kift* is not (Ganong, 1980). Ambiguous tokens may also be interpreted to maximize neighborhood density—e.g., ambiguous [k]~[g] is more likely to be interpreted as /g/ in the context /___ice, since the non-word [gais] has a higher neighborhood density than [kais] does (Newman, Sawusch & Luce, 1997). Finally, ambiguous tokens may be interpreted in a way that maximizes the phonotactic probability of the resulting word—e.g., ambiguous [r]~[l] is more likely to be identified as /r/ in the context /t___, and as /l/ in the context /s___, since **tl* and **sr* are phonotactically illegal in English (Massaro & Cohen, 1983; Moreton, 2002). The phonotactic expectations that drive perceptual biases can be complex and non-local. For example, Coetzee (2008) shows that English speakers are biased to interpret ambiguous [t]~[k] tokens as [t] in the context /skV___, in order to avoid a phonotactically dispreferred *skVk* sequence. Certainly, some of the contexts discovered by the learning algorithm as favoring [np] are likely to be too complex and non-local to create such perceptual biases, but it is conceivable that at least some of these emergent generalizations could guide perception of further tokens, cementing and strengthening themselves over time.

The upshot of this section, then, is that emergent coincidental subregularities may be a more common threat than one might imagine. Although there was just one systematic factor shaping these languages (a following [p] tends to favor [m] over [n]), other patterns may emerge “accidentally”. This is especially likely in cases where there are relatively small numbers of forms involved, but this is often the case in language data—especially in the sets of words available to children. A powerful inductive learning algorithm which is not biased against such generalizations may find such patterns, and extend them through feedback and reinforcement. This mechanism has the potential to predict a pathway to unnatural phonologies: languages in the midst of a phonetically motivated change get diverted by minor subregularities, stopping the natural change prematurely and ending in an unnatural distribution. Cases best described in these terms do not appear to be widespread in the literature.³ We may provisionally

³There is, of course, a possibility that such cases do exist, but that they have not been recognized as such because this mechanism of change has not generally been considered. It is hoped that the present work may serve as inspiration to uncover such cases, if they do exist.

conclude that the learning mechanism is biased against finding such contexts.

3. A role for grammatical biases

The desired effect of a learning bias is to prevent the learner from taking seriously contexts such as the one in (4) above. A direct way of accomplishing this would be to build in a prior bias for phonetically natural constraints such as *[np], and against unnatural constraints such as the one in (4). However, the current results do not necessarily motivate such a move, since it is entirely possible that more generic biases could also accomplish the same goal. In this section, we explore the role of three other types of bias: simplicity, locality, and sensible natural classes.

3.1. A simplicity bias

The constraint in (4) is clearly much more complicated than the more natural *[np] constraint. Thus, one obvious line of attack is to penalize constraints that are too complex. This approach would help in eliminating many undesirable constraints, but also has some drawbacks and limitations. For one thing, legitimate phonological contexts may also be somewhat complicated, as in the *skVk example cited above. Therefore, the simplicity bias cannot be so severe that it completely disallows complex environments. However, if we allow complex environments, then the model will find and make use of them, re-introducing at least some of the unwanted unnatural constraints. The difficulties of using structural complexity alone as evaluation metric are well known; see, for example, Chomsky & Halle (1968, chap. 9). Preferences for simpler patterns have indeed been documented in the literature (e.g., Seidl & Buckley 2005), and inductive models may well benefit from incorporating a simplicity bias. However, it seems unlikely that a bias against complex conditioning contexts could completely avoid the unwanted contexts observed here.

3.2. A locality bias

A more promising bias is to avoid constraints that refer to contexts not immediately adjacent to the target. This can be implemented as a learning bias: when comparing words, focus preferentially on the material that is immediately adjacent to the change, ignoring material that is farther away. It may also be implemented as an evaluation bias in which a numerical penalty is assessed to referring to non-local material, with the penalty perhaps increasing in proportion to the distance from the change. The importance of locality can be demonstrated with an example from English past tense allomorphy. The regular English past suffix has three allomorphs ([t], [d], [əd]), with the choice almost completely determined by the preceding segment: [əd] after [t, d] (e.g., *want*[əd], *need*[əd]), [t] after voiceless segments (e.g., *laugh*[t], *miss*[t], *hope*[t]), and [d] elsewhere (e.g., *plan*[d], *pull*[d], *amuse*[d]). In addition, there are cases of “hyperregularity” (islands of reliability for regular marking) conditioned by the preceding segment: as it turns out, all verbs ending in voiceless fricatives are regular, and speakers appear to notice this fact, giving such verbs regular inflection at greater than average rates (Albright & Hayes, 2003). Although they have received less attention in the literature, there are also cases of hyperregularity conditioned by non-adjacent segments. For example, verbs beginning with voiced fricatives are also all regular in English.⁴ However, speakers appear not to notice this context. Prasada & Pinker (1993) collected ratings of regular and irregular past tense forms for a variety of nonce verbs, including one pair that differed in the voicing of initial fricatives. The results in (6) suggest that initial voiced fricatives are not preferentially associated with regular past tense marking—in fact, quite the opposite.

- (6) Initial voiced fricatives do not condition regular past tense marking (Prasada & Pinker, 1993)

Novel verb	Regular rating	V-change rating
<i>foa</i> [fɔv]	5.9	5.0
<i>voa</i> [vɔv]	2.8	5.4

⁴There is a historical explanation for this fact: Old English did not have a voicing contrast among fricatives, and all fricatives were voiceless word-initially. Most verbs beginning in *v* and *z* are borrowed from French (*view*, *vow*) or are onomatopoeic (*zap*, *zoom*), and are thus regular.

Although this result is based on a comparison between just two forms and is thus suggestive at best, we take it as an indication that non-local conditioning contexts may be weaker than local ones. In the present case, a bias against non-local conditioning contexts may attenuate the advantage of constraints such as the one in (4), making them weaker and less likely to be enforced.

3.3. *Sensible natural classes*

One last salient property of the constraint in (4) is that it makes use of negatively defined natural classes, such as “non-nasal and non-lateral” ([–nasal,–lateral]). In fact, a large proportion of the 678 unnatural generalizations reported above make use of such classes, which refer to sets of segments that share nothing in common except the lack of particular features. Such classes are a well-known problem in feature theory, since in many cases they are never found to play a role in defining phonotactic distributions or conditioning alternations. For this reason, they are often eliminated by use of privative features, which have only positive values and no negative values. In the present case, defining [nasal] and [lateral] as privative features would prevent the model from referring to classes such as [–nasal,–lateral]. As a result, the hypothesis space for possible constraints would be much smaller, and the number of possible unnatural generalizations would shrink considerably.

In order to test the effect of privative specifications, I reran the learning simulations of the 1,000 VNpV languages, this time using a feature system that incorporated privative specifications for features such as nasal and lateral, as well as for major places of articulation. The result was that all of the unnatural generalizations reported above were eliminated. This solution has a potential advantage over simplicity and locality biases, since we know that phonological contexts may be somewhat complex and non-local (precluding too steep a simplicity or locality bias), but it appears that contexts such as [–nasal,–lateral] are rarely if ever relevant in defining phonological distributions. Constraining the model to make use of only phonetically sensible natural classes is not the same thing as building in prior knowledge of phonetically sensible markedness constraints, but it does represent a form of substantive phonological knowledge. The results here demonstrate the usefulness of such knowledge, even if it is not the only possible solution in the present case.

4. Conclusion

In this paper, I have tested the idea that a power inductive learner, fed an input of training data that is shaped by misperception, may reinforce observed asymmetries and help eliminate typologically marked contrasts. Learning trials using an inductive learner reveal that accidentally true patterns may arise quite often, creating an unexpected pathway by which natural change may get sidetracked and unnatural languages may arise. This mechanism has the potential to undermine the expected effect of “transmission bias”. Such accidental correlations may be safely ignored by the learner if we build in certain grammatical biases, such as simplicity, locality, and phonetically sensible natural classes. So far, the most significant improvements appear to come from locality and phonetically sensible classes, though much work remains to be done testing the viability of such biases in modeling data from a variety of languages and tasks.

More generally, the results of the current study suggest that caution is needed when predicting that certain patterns are unlikely to arise, when we are dealing with complex systems and limited data. In fact, accidentally true and entirely unnatural patterns are all around us, but only a small number of these appear to be noticed and grammaticalized. Much empirical work is needed, both in looking for such patterns in actual languages, and in testing the extent to which speakers notice or ignore them. The fact that such patterns appear to be very rare typologically is a preliminary indication that the minimally biased model described above pays too much attention to them, and overpredicts their rate of occurrence. Thus, although isolating the correct mix of biases is left for future research, it seems quite certain that some sort of bias must be built into the model in order to predict the observed typology.

References

- Albright, Adam (2002). Islands of reliability for regular morphology: Evidence from Italian. *Language* 78:4, pp. 684–709.
- Albright, Adam & Bruce Hayes (2003). Rules vs. analogy in English past tenses: A computational/experimental study. *Cognition* 90, pp. 119–161.
- Albright, Adam & Bruce Hayes (2006). Modeling productivity with the Gradual Learning Algorithm: The problem of accidentally exceptionless generalizations. Fanselow, Gisbert, Caroline Féry, Ralf Vogel & Matthias Schlesewsky (eds.), *Gradience in Grammar: Generative Perspectives*, Oxford University Press, pp. 185–204.
- Becker, Michael, Nihan Ketrez & Andrew Nevins (2007). The surfeit of the stimulus: Analytic biases filter the statistics of Turkish voicing.
- Berent, Iris, Donca Steriade, Tracy Lennertz & Vered Vaknin (2007). What we know about what we have never heard: Evidence from perceptual illusions. *Cognition* 104, pp. 591–630.
- Blevins, Juliette (2004). *Evolutionary Phonology: The Emergence of Sound Patterns*. Cambridge University Press.
- Blevins, Juliette & Andrew Garrett (2004). The evolution of metathesis. Hayes, Bruce, Robert Kirchner & Donca Steriade (eds.), *Phonetically based phonology*, Cambridge University Press.
- Bybee, Joan (2001). *Phonology and language use*. Cambridge University Press, Cambridge.
- Chomsky, Noam & Morris Halle (1968). *The sound pattern of English*. Harper and Row, New York.
- Coetzee, Andries (2008). Grammaticality and ungrammaticality in phonology. *Language* 84, pp. 218–257.
- Ernestus, Mirjam & R. Harald Baayen (2003). Predicting the unpredictable: Interpreting neutralized segments in Dutch. *Language* 79, pp. 5–38.
- Finley, Sara & William Badecker (in press). Towards a substantively biased theory of learning. *BLS* 33.
- Frisch, Stefan A., Janet B. Pierrehumbert & Michael B. Broe (2004). Similarity avoidance and the ocp. *Natural Language & Linguistic Theory* 22:1, pp. 179–228.
- Ganong, William F. (1980). Phonetic categorization in auditory word perception. *Journal of Experimental Psychology: Human Perception and Performance* 6, pp. 110–125.
- Harris, Alice C. (2008). On the explanation of typologically unusual structures. Good, Jeff (ed.), *Linguistic Universals and Linguistic Change*, Oxford University Press, pp. 54–76.
- Hayes, Bruce (1999). Phonetically-driven phonology: The role of Optimality Theory and inductive grounding. Darnell, Michael, Edith Moravcsik, Michael Noonan, Frederick Newmeyer & Kathleen Wheatly (eds.), *Functionalism and Formalism in Linguistics, Volume I: General Papers*, John Benjamins, Amsterdam, pp. 243–285.
- Hayes, Bruce & Donca Steriade (2004). The phonetic basis of phonological markedness. Hayes, Bruce, Robert Kirchner & Donca Steriade (eds.), *Phonetically based phonology*, Cambridge University Press, Cambridge, pp. 1–33.
- Hura, Susan, Björn Lindblom & Randy Diehl (1992). On the role of perception in shaping phonological assimilation rules. *Language and Speech* 35, pp. 59–72.
- Kawahara, Shigeto (2006). Mimetic gemination in Japanese: A challenge for Evolutionary Phonology. *Theoretical Linguistics* 32–3, pp. 411–424.
- Kochetov, Alexei (2004). Positional markedness as a by-product of the learning situation. Nowak, Pawel M., Corey Yoquelet & David Mortensen (eds.), *Proceedings of the 29th meeting of Berkeley Linguistics Society*, University of California, Berkeley, pp. 545–556.
- Kruskal, Joseph B. (1983). An overview of sequence comparison. *Time Warps, String Edits, and Macromolecules: The Theory and Practice of Sequence Comparison*, Addison-Wesley, Reading, MA, pp. 1–44.
- Massaro, D. W. & M. Cohen (1983). Phonological context in speech perception. *Perception and Psychophysics* 34, pp. 338–348.
- Moreton, E. (2002). Structural constraints in the perception of English stop-sonorant clusters. *Cognition* 84, pp. 55–71.
- Moreton, Elliott (2007). Learning bias as a factor in phonological typology. Paper presented at the 26th meeting of the West Coast Conference on Formal Linguistics, Berkeley, California, April 27–29.
- Morley, Rebecca (2008). Likely input and the reasonable learner. Johns Hopkins ms.
- Newman, R. S., J. R. Sawusch & P. A. Luce (1997). Lexical neighborhood effects in phonetic processing. *Journal of Experimental Psychology: Human Perception and Performance* 23, pp. 873–889.
- Ohala, John (1981). The listener as a source of sound change. *Chicago Linguistic Society*, vol. 17, pp. 178–203.
- Peperkamp, Sharon, Rozenn Le Calvez, Jean-Pierre Nadald & Emmanuel Dupoux (2006). The acquisition of allophonic rules: Statistical learning with linguistic constraints. *Cognition* 101, pp. B31–B41.
- Pierrehumbert, Janet B. (2001). Exemplar dynamics: Word frequency, lenition and contrast. Bybee, Joan & Paul Hopper (eds.), *Frequency and the emergence of linguistic structure*, John Benjamins, pp. 137–158.
- Pierrehumbert, Janet B. (2006). The statistical basis of an unnatural alternation. Goldstein, L. M., D. H. Whalen & C. T. Best (eds.), *Laboratory Phonology VIII: Varieties of Phonological Competence*, Mouton de Gruyter, Berlin.
- Prasada, Sandeep & Steven Pinker (1993). Generalization of regular and irregular morphological patterns. *Language and Cognitive Processes* 8, pp. 1–56.

- Prince, Alan & Paul Smolensky (2004). *Optimality Theory: Constraint Interaction in Generative Grammar*. Blackwell Publishing.
- Pycha, Anne, Pawel Nowak, Eurie Shin & Ryan Shosted (2003). Phonological rule-learning and its implications for a theory of vowel harmony. Garding, G. & M. Tsujimura (eds.), *WCCFL 22 Proceedings*, Cascadilla Press, Somerville, MA, pp. 423–435.
- Seidl, Amanda & Eugene Buckley (2005). On the learning of arbitrary phonological rules. *Language Learning and Development* 1, pp. 289–316.
- Steriade, Donca (2001). Directional asymmetries in place assimilation: a perceptual account. Hume, Elizabeth & Keith Johnson (eds.), *Perception in Phonology*, Academic Press, pp. 219–250.
- Wedel, Andrew (2007). Feedback and regularity in the lexicon. *Phonology* 24, pp. 147–185.
- Wilson, Colin (2003). Experimental investigation of phonological naturalness. Garding, G. & M. Tsujimura (eds.), *WCCFL 22 Proceedings*, Cascadilla Press, Somerville, MA, pp. 533–546.
- Wilson, Colin (2006). Learning phonology with substantive bias: An experimental and computational study of velar palatalization. *Cognitive Science: A Multidisciplinary Journal* 30, pp. 945–982.
- Winters, Stephen (2003). *Empirical Investigations into the Perceptual and Articulatory Origins of Cross-Linguistic Asymmetries in Place Assimilation*. Ph.D. thesis, The Ohio State University.
- Zhang, Jie & Yuwen Lai (2006). Testing the role of phonetic naturalness in Mandarin tone sandhi. *Kansas Working Papers in Linguistics* 28, pp. 65–126.
- Zuraw, Kie (2000). *Patterned Exceptions in Phonology*. Ph.D. thesis, UCLA.
- Zuraw, Kie (2007). The role of phonetic knowledge in phonological patterning: Corpus and survey evidence from Tagalog. *Language* 83, pp. 277–316.