To appear in:
Proceedings of the 29th meeting of the Berkeley Linguistic Society, 2003

Positional markedness as a by-product of the learning situation[*]

ALEXEI KOCHETOV
Simon Fraser University and Haskins Laboratories

0. Introduction

This paper presents a simulation where a "learner" attempts to acquire a positionally unrestricted phonological contrast (plain vs. palatalized labial stops in a variety of contexts) without any a priori knowledge of the relevant positional hierarchies. The results show that an acquisition model that includes human-like perceptual limitations can give rise to cross-linguistically attested phonotactic grammars without positing this a priori knowledge.

1. Positional scales and learnability

Scales encoding relative markedness of environments, referred to as "positional markedness scales," play an important role in Optimality Theory (Prince & Smolensky 1993). For instance, the contrast between voiceless and voiced obstruents is known to be commonly neutralized in syllable coda position (Lombardi 1995), or, under a different analysis, in the absence of a following sonorant segment (which contains acoustic cues to the contrast: Steriade 1997). These facts can be seen as reflecting universal positional markedness scales that refer to syllable positions or phonetically-cued environments. The scales can be represented grammatically as hierarchies of positional markedness constraints (or faithfulness constraints: Beckman 1997). It is important to note that these constraint hierarchies are assumed to be *harmonic*, or universally fixed, and thus innate by definition (Prince & Smolensky 1993).

A traditional OT learner constructs a phonological grammar by ranking markedness and faithfulness constraints based on positive evidence from the lexicon s/he has acquired (Tesar & Smolensky 1998, Tesar 2000). UG supplies the learner not only with the relevant universal constraints, but also with positional (and context-free) markedness hierarchies. The question of whether the lexical items the learner is exposed to are always perfectly produced and perfectly

---

[*] Many thanks to Keren Rice, Louis Goldstein, Elan Dresher, Maria Babyonyshev, Edward Flemming, Elizabeth Hume-O'Haire, Keith Johnson, Dasha Kavitskaya, Hank Rogers, Joseph Schallert, Ron Smyth, and Donca Steriade, among others. The research was supported by fellowships from SSHRC, Canada. All errors are my own.

recoverable is hardly addressed in the OT literature on learnability (but has been raised elsewhere: Ohala 1981 et seq.).

In this paper I take a different approach to positional markedness scales, arguing that it is unnecessary to posit a priori knowledge of the scales (with or without direct reference to phonetic factors) because the attested patterns of neutralization emerge naturally from speech production and perception during language acquisition without this knowledge. The case study explored here is a positional markedness scale that reflects cross-linguistic facts of neutralization of the contrast between plain and palatalized labials.

2. A scale of distribution of the plain/palatalized contrast

A survey of over 20 languages and dialects with contrastive palatalization (Kochetov 2002) shows the following patterns: the contrast between plain (or velarized) and palatalized labials, such as /p/ vs. /p^j/, is most often restricted to the prevocalic position and is neutralized elsewhere (Type 1: Table 1); only a small number of languages allow the contrast both prevocalically and word-finally (Type 2); only one language extends the /p/ vs. /p^j/ distinction to the context before plain consonants (Type 3). A full contrast (Type 0), that is a pattern in which the contrast is expressed in all environments, appears to be unattested.[1] The distribution is also dependent on the quality of the following vowels and the place of the following consonants, which are not considered here. Note that the most common output of neutralization is a plain consonant.

Table 1. A summary of distribution of the plain/palatalized contrast in labials

Patterns	_V pa vs. p^{j}a	_# ap vs. ap^j	_C apta vs. ap^{j}ta	_C^j apt^{j}a vs. ap^{j}t^{j}a	Comments
Type 1	yes	no	no	no	common
Type 2	yes	yes	no	no	rare
Type 3	yes	yes	yes	no	very rare
Type 0	yes	yes	yes	yes	unattested

The observed asymmetries can be represented by a scale of environments (1a) arranged from the least marked (before a vowel) to the most marked (before a palatalized consonant). Assuming that this scale is universal, that is, that it holds in all languages that maintain the plain-palatalized contrast in labials, a standard approach would be to posit a fixed hierarchy of markedness constraints that prohibit the plain/palatalized contrast (*[αpal]) in specific contexts, as in (1b) (cf., the constraints on *[αvoice] in Steriade 1997). The differences between the grammars of different language will then result from different rankings of the

---

[1] Some examples of the languages surveyed include Standard Bulgarian, Lithuanian, Nenets (Type 1); Russian, Irish (Type 2); and the Nova Nadezhda dialect of Bulgarian (Type 3) (see Kochetov 2002 for additional detail as well as references).

faithfulness constraint Ident[αpal] against the hierarchy of positional markedness constraints.

(1) a. __V > __# > __C > __C^j
b. *[αpal]/_C^j » *[αpal]/_C » *[αpal]/_# » *[αpal]/_V

The question, however, is whether the scale in (1) and the corresponding constraint rankings can arise without being pre-specified in UG. This possibility is explored by means of the learning simulation described below.

3. The learning situation

The approach explored here builds in part on Hayes' (2001) model of a "pure phonotactic learner." This learner, provided with a set of well-formed words in a given language and universal OT constraints, constructs a grammar (a ranking of constraints) that can evaluate well-formed and ill-formed phonotactic sequences (without reference, at this stage, to their morphological composition and alternations). At the same time, the current analysis represents a departure from other learnability work within OT. First, it assumes a more realistic speaker/listener who has inherent limitations on speech production and perception. Second, our learner is *not* provided with implicit knowledge of phonological markedness, in particular, the positional markedness scale. In other words, s/he is not aware of whether the contrast is more marked (or harder to produce/perceive) in some environments than in others (cf., Ohala 1981, Hale & Reiss 2000, Hume & Johnson 2001; cf. Mielke, this volume).[2] The current approach also relies crucially on the concept of *self-organization*, or spontaneous emergence of order, that is characteristic of many natural and artificial dynamical systems (see, for example, Kauffman 1995, Langton 1995; cf. Wedel, this volume). Specifically, in our case, the self-organization approach holds that positional markedness scales – high-level phonological structure – is a by-product of "blind" low-level interactions between speakers and listeners during the learning process.

This hypothesis with respect to the scale in (1) is investigated using a computer simulation of the interactions of agents, or simple autonomous entities. Agent-based programming provides explicit ways of testing relative importance of multiple factors, and it has been used recently to explore a variety of emergent phonological phenomena (e.g., Browman & Goldstein 1999, de Boer 2000, Liberman to appear). The current, rather simplistic, *Matlab* simulation involves two agents, an adult agent and a learning agent; both consist of components modeling speech production, perception, the lexicon, and grammar. The focus here is on the learner's ability to perceive lexical items produced by the adult, to build a lexicon, and to construct a grammar. My assumptions about the target grammar and the adult-learner interactions are presented in the following sections.

---

[2] The learner, however, may infer some knowledge of relative markedness (e.g., perceptibility: Steriade 2001; or articulatory complexity: Hayes 1999) from sensory/motor experience.

3.1. The hypothetical language

Let's assume a hypothetical language, Language X, that has the consonant inventory {p, p^j, t, t^j} and a vowel {a}. The lexical items of the language are presented in (2a); they consist of mono-morphemic words with a contrast between plain and palatalized labials in four environments: word-initially, word-finally, and before plain and palatalized consonants. Assume also that the knowledge of well-formed sequences of segments in this language, the phonotactic grammar, is encoded by ranking the Optimality Theory-type markedness and faithfulness constraints presented in Section 2. The grammar, that is, the ranking of these constraints characteristic of Language X, is shown in (2b). This grammar, unattested cross-linguistically (Type 0 in Table 1) will be the target grammar for our hypothetical learner.

(2) a. Lexicon: {pa, p^ja, ap, ap^j, apta, ap^jta, apt^ja, ap^jt^ja}
b. Grammar: Ident[αpal] » *[αpal]/_C^j, *[αpal]/_C, *[αpal]/_#, [αpal]/_V

3.2. The adult agent

The adult agent has full knowledge of the lexicon and the grammar. His/her goal in this simulation is to "produce" lexical items. This production is not directly simulated; instead, it is assumed to correspond to the actual production of the relevant consonants *at word boundaries* in nonsense utterances (Table 2). These data were obtained using an articulatory magnetometer from four speakers of Russian. It should be noted that the contrast between /p/ and /p^j/ is present in all contexts (_V, _#, _C, _C^j), although there are some gradient differences in the magnitude of tongue body raising/backing gestures and their timing relative to the lips (Kochetov, to appear). What is important here is that the adult agent's production is imperfect, having the kinds of limitations typical of a human speaker.

Table 2. Adult agent's "production"

Environments	Lexical items, Language X	Russian nonwords
(#)_V	pa, p^ja	ta [p]apy, ta [p^j]apy
(V)_#	ap, ap^j	ta[p] apy, ta[p^j] apy
(V)_C	apta, ap^jta	ta[p] tapy, ta[p^j] tapy
(V)_C^j	apt^ja, ap^jt^ja	ta[p] t^japy, ta[p^j] t^japy

3.3. The learning agent

Unlike the adult agent, the learning agent begins with an empty lexicon (3a), and the grammar (3b), where the markedness constraints are ranked above the faithfulness constraint (cf., Hayes 2001). The goals of the learning agent are thus (i) to recover presented items from an acoustic signal, (ii) to posit representations of these items based on a given number of tokens per item, and (iii) to rank constraints in the grammar based on generalizations over stored lexical items.

(3) a. Lexicon: {Ø}
b. Grammar: *[αpal]/_C^j, *[αpal]/_C, *[αpal]/_#, *[αpal]/_V » Ident[αpal]

Perception

The agent's ability to recover items from the signal is limited; it is based on the probability with which humans tend to perceive the contrast /p/ vs. /p^j/ (Kochetov 2002).

The perceptibility scale used here (Figure 1a) reflects the likelihood of each segment being perceived as "palatalized" [1] or "plain" [0] in a given environment *at word boundaries* (where no phonotactic restrictions apply in Russian). It represents the average perception of the stimuli given in Table 2 under a number of experimental conditions: by native listeners without noise (N = 20), by native listeners with noise (N = 20), and by non-native listeners without noise (Japanese, N = 10).[3] A scale of the perceptual similarity of /p/ and /p^j/ in different phonological contexts (Figure 1b) was computed from confusion matrices (based on correct responses and false alarms; see Johnson 2003). Note that the perceptual scales are assumed here to reflect universal performance biases in auditory perception that are external to the phonological grammar proper (cf., Hume & Johnson 2001).

As we can see in Figure 1, both the "plain" or "palatalized" responses and the degree of similarity between the segments are substantially affected by environments: there is an almost 100% correct identification rate of /p/ and /p^j/ in the prevocalic context, but a high rate of confusion word-finally and before consonants (especially for /p^j/).[4] The positional effect is striking, especially given the fact that all these word-boundary contexts are phonotactically possible both for native and non-native listeners (the Japanese listeners interpreted coda [p] and [p^j] as syllables /pu/ and /pi/).

Figure 1. Mean perception of /p/ and /p^j/ as "plain" [0] or "palatalized" [1] in four contexts at word boundaries (a); mean perceptual similarity between /p/ and /p^j/ in four contexts (b); "0" is "different," "1" is "same."

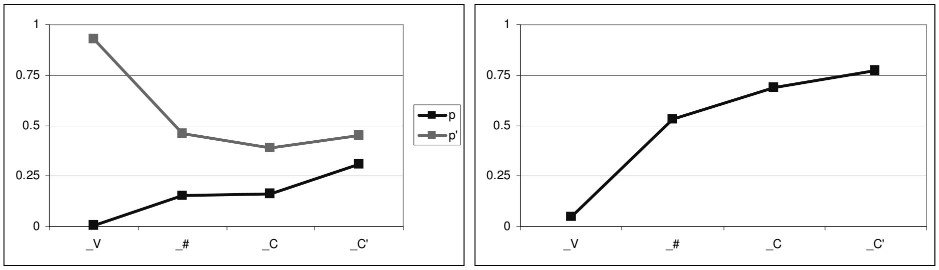
a. b.

---

[3] Non-native listeners were used as an additional control for language-specific biases.

[4] Articulatory and acoustic factors that may have contributed to these results are discussed in Kochetov (2002).

To model the likelihood of the consonants being perceived as “plain” or “palatalized,” the simulation employs a random number generator. For instance, if a random number from 0 to 1 is less than 0.93 (see Figure1a) the item {p^ja} is “perceived” by the listener/learner as {p^ja} (“palatalized”). If the number is greater than 0.93, the item is “perceived” as {pa} (“plain”). The perceived tokens – whether correctly identified or confused with other forms – are stored in memory. In each of the tokens the consonant is specified as “palatalized” by assigning it the value [1], or “plain,” by assigning it the value [0].

Lexicon and grammar

Based on the stored tokens, the learner posits a single form for a given lexical item, that is, its representation. The representation of the form is determined based on the following labeling algorithm: If the mean for all the stored tokens for a given item has a score of less than 0.45 (i.e., most tokens are “plain”), then the item is judged to have a non-palatalized consonant, [0]. If the score is higher than 0.55 (i.e., most tokens are “palatalized”), the item is judged to have a palatalized consonant, [1]. If the item falls within the range 0.45-0.55, it is randomly assigned one of the two labels. For example, if 56 out of 100 tokens of {ap^j} are perceived as “plain” (an average of 0.56) and 44 tokens are perceived as “palatalized” (an average of 0.44), the form is labeled, or “saved” in the lexicon, as {ap^j}. A lexicon is represented as a matrix of ones and zeros: for example, [0 0 0 1; 1 0 0 1] represents a lexicon containing the items {pa, ap, apta, ap^jt^ja; p^ja, ap, apta, ap^jt^ja}. These initial lexical representations are updated based on the learner’s perception of his/her own production of pairs of items (e.g., {apt^ja} and {ap^jt^ja}). If the pairs of items are perceived as being similar, their representations are revised (e.g., to either {apt^ja} or {ap^jt^ja}). The two pairs are considered similar if a generated random number is below the perceived similarity score for this pair (Figure 1b).

Based on the lexicon acquired, the learner constructs the grammar by ranking the constraints. Recall that the initial state of the learning agent is the ranking in (3b): the faithfulness constraint is ranked at the bottom of the hierarchy and the positional markedness constraints are not pre-ranked with respect to each other. This is important, because the constructed grammar becomes crucially dependent on the items the learner stores in the lexicon. In the case of perfect perception the agent’s lexicon and grammar would be an exact reflection of the target lexicon and grammar. If the learner’s perception is completely random, we may expect a wide range of possible lexicons and grammars (16 logically possible permutations of the constraints in (3b)).

In sum, the processing of the presented items by the learner thus consists of recovering these items and making several types of simple generalizations: generalizations over tokens of an item (lexical representations), over pairs of lexical items (lexical contrast), and over the lexicon in general (the grammar). Although the learning situation described here is a much simplified model of a real-life acquisition path, the simplifications will allow us to focus on the phenomenon under investigation, as well as to keep the simulation computationally tractable.

Having presented the mechanism of the simulations, I turn to the following question: Is our "imperfect" learner capable of developing the target lexicon and constructing the target grammar, and if not, what are the resulting lexicons and grammars?

4. Results

In this section I describe the results of the simulation representing a sample run based on 100 tokens per item and 100 iterations. This means that our learner is presented each time with 100 tokens of 8 lexical items from Lexicon 0. The learner recovers these tokens and posits 8 representations. The representations are updated after 100 productions of pairs of items and these updated representations then constitute the learner's lexicon. Based on the items in the lexicon, the learner ranks 5 constraints, constructing the grammar. This process is repeated 100 times (with the same input, Lexicon 0), that is, one run of the simulation produces 100 grammars. After reviewing the results, presented separately for the lexicons (Section 4.1) and the grammars (Section 4.2), I address the issue of sensitivity of the results to parameter settings (Section 4.3) and relevance of other factors (Section 4.4).

4.1. Building the lexicon

Table 3 summarizes the results of a sample run in terms of major lexicon types; any deviations from the target lexicon are shaded. Lexicon 0 is the target lexicon that that contains 8 phonologically distinct lexical items. Note that this lexicon is not once replicated by the learner (0 occurrences). Lexicon 1 shows a high degree of homophony, with {ap, apta, apt^{j}a or apt^{j}a} corresponding to the original pairs {ap vs. ap^j, apta vs. ap^{j}ta, apt^{j}a vs. ap^{j}t^{j}a}. This lexicon is by far the most common outcome: it was generated 74 out of 100 times.[5] Lexicon 2, which differs from the previous lexicon in preserving the original distinction {ap vs. ap^j} was produced by the learner only 12 times. Lexicon 3 is listed here because it is minimally different from the target Lexicon (by the lack of the contrast {apt^{j}a vs. ap^{j}t^{j}a}); this lexicon occurs in the output only once. It is thus as likely as (or even less likely than) some other generated lexicons, grouped under the title "Other lexicons". In the current sample these include three lexicons that collapse the distinctions between {ap vs. ap^j} or {ap^{j}ta vs. ap^{j}t^{j}a} in favor of the palatalized forms {ap^j ap^{j}ta}.[6] Note that all of the resulting lexicons keep the distinctions between the items {pa} and {p^{j}a}.

---

[5] The results for the last environment (_C^j) are combined here; there were 57 occurrences of {apt^{j}a} and 17 occurrences of {ap^{j}t^{j}a}.

[6] This result is due to the learner's random selection of one of the relatively similar forms (see Section 3.3).

It should be mentioned that additional runs under the same conditions showed very similar results: the absence of Lexicon 0, and relatively high frequency of Lexicons 1 and 2. Overall, the degree of variation among the lexicons generated is quite limited, with certain types of lexicons clearly preferred over others.

The results of the simulation are not surprising given the learner's human-like perceptual limitations and biases: The learning agent tends to perceive the item {ap^{j}} as [ap] almost as often as [ap^{j}]; the same, however, does not hold for the item {ap} (see Figure 1). The pairs {apta} vs. {ap^{j}ta} and {apt^{j}a} vs. {ap^{j}t^{j}a} have an even higher error rate and perceived similarity. Generalizations over tokens and pairs of items help the learner select one form or the other, but they do not eliminate the strong biases towards certain forms (e.g., {ap} or {apta}).

Table 3. Resulting lexicons: A sample run based on 100 tokens (100 iterations)

Lexicon 0		Lexicon 1		Lexicon 2		Lexicon 3		Other lexicons	
0 occurrences		74 occurrences		12 occurrences		1 occurrence		13 occurrences	
pa	p^{j}a	pa	p^{j}a	pa	p^{j}a	pa	p^{j}a	pa	p^{j}a
ap	ap^{j}	ap		ap	ap^{j}	ap	ap^{j}	ap or ap^{j}	
apta	ap^{j}ta	apta		apta		apta	ap^{j}ta	apta or ap^{j}ta	
apt^{j}a	ap^{j}t^{j}a	apt^{j}a or ap^{j}t^{j}a		apt^{j}a or ap^{j}t^{j}a		apt^{j}a		apt^{j}a or ap^{j}t^{j}a	

4.2. Constructing the grammar

Having posited the lexical forms, the learner constructs a subset of the phonotactic grammar, a set of generalizations about how segments can be combined in the language. This is done by demoting positional markedness constraints based on evidence from the lexicon (e.g., Hayes 2001).

Unsurprisingly, our results reveal zero occurrences of the target grammar, as shown in Table 4. Recall that, in this least restrictive grammar, the faithfulness constraint Ident[αpal] is ranked above all the positional markedness constraints. The set of grammars the learner does arrive at is very limited. The most common outcome by far is Grammar 1, which neutralizes the contrast in all contexts except the prevocalic one. Note that, for simplicity, the constraints used do not specify whether the outcome of neutralization is a plain or a palatalized segment (but see Kochetov 2002). Thus, Grammar 1 can be constructed based on any lexicon that preserves the original distinction among {pa} vs. {p^{j}a}, but not between other items (Lexicon 1 and the lexicons listed in the last column). Grammar 2, a consequence of positing Lexicon 2, is much less common. Finally, given the unlikely Lexicon 3, the corresponding grammar (Grammar 3) is also a rare outcome. Additional runs under the same conditions replicate the strong bias towards Grammars 1 and 2.

Note that we do not find in the results a number of other logically possible rankings (in fact, 12 out of 16 logically possible rankings). These include, for instance, grammars that disallow the contrast word-initially while maintaining it word-finally or preconsonantally. Given the probabilistic nature of the learner's perception, there is no categorical prohibition against these rankings (as well as the target grammar) — these grammars are theoretically possible, yet very unlikely due to the learner's perceptual abilities. It should be noted, however, that other factors (discussed below) can override the substantive biases.

Table 4. Resulting grammars: A sample run based on 100 tokens (100 iterations)

Grammar 0	Grammar 1	Grammar 2	Grammar 3	Other grammars
0 occurrences	87 occurrences	12 occurrences	1 occurrence	0 occurrences
Ident[αpal] » *[αpal]/_C^j, *[αpal]/_C, *[αpal]/_#, *[αpal]/_V	*[αpal]/_C^j *[αpal]/_C, *[αpal]/_#} » **Ident[αpal]** » *[αpal]/_V	*[αpal]/_C^j *[αpal]/_C} » **Ident[αpal]** » *[αpal]/_#, *[αpal]/_V	*[αpal]/_C^j » **Ident[αpal]** » *[αpal]/_C, *[αpal]/_#, *[αpal]/_V	…

The most significant outcome of our simulation is that the results reflect in general the cross-linguistically attested patterns of neutralization of the plain-palatalized contrast (see Table 1). Thus the most commonly generated grammar corresponds to the most commonly attested pattern of neutralization of the plain-palatalized contrast in labials (Type 1). Grammar 2 is a less common outcome; it is also less common cross-linguistically (Type 2). Finally, the very unlikely outcome, Grammar 3, is attested in only a single language. The simulation thus makes correct predictions about the range of cross-linguistic variation with respect to the neutralization of the contrast, as well as to the relative frequency of each particular pattern. Crucially, to the analysis that the positional markedness scale (_V > _# > _C > _C^j) has been achieved without or independently of the learner having prior knowledge of it.

4.3. Parameter settings

The question we need to address next is to what extent the results discussed in the previous sections are influenced by the parameters specified in the current simulation. Recall that our learner was presented with 100 tokens of each lexical item. Lexical representations were posited based on the following, quite deterministic, labeling algorithm: *if* the mean score $x \in$ 0-0.45 *then* "plain", *if* $x \in$ 0.55-1 *then* "palatalized", *if* $x \in$ 0.45-0.55 *then* either of the two (see Section 3.3). Subsequent manipulations of these parameters have shown them to affect the results to a certain degree, but within important limits.

Figure 2 plots the mean contrast between /p/ and /p^j/ in each environment (based on 100 iterations) resulting from four different settings: either 10 or 100 tokens categorized by the algorithm described above ("10 strict" and "100 strict"); either 10 or 100 tokens categorized by a less deterministic labeling algorithm ("10 lax" and "100 lax"): *if* the mean score x ∈ 0-0.25 *then* "plain", *if* x ∈ 0.75-1 *then* "palatalized", *if* x ∈ 0.25-0.75 *then* either of the two. Note that under all four settings the simulation results in grammars all of which maintain the contrast prevocalically (with 3 exceptions under "10 lax"). The differences are primarily in the degree of variability (maintained or neutralized) in the less favorable environments (_#, _C, and C^j); with all of the settings, however, the contrast is supported in a relatively small number of grammars. Overall, fewer tokens and less deterministic labeling procedures lead to somewhat higher variability in the output.

Figure 2. Resulting contrasts between /p/ and /p^j/ in 4 environments based on 10 or 100 tokens and "strict" or "lax" labeling algorithms; the mean for 100 iterations; 1 = "contrast", 0 = "no contrast".

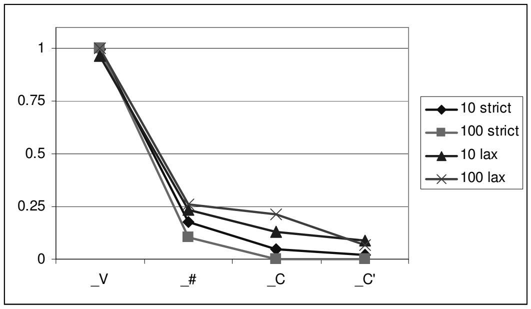

4.4. Limitations and additional factors

A number of questions related to the results of the simulation require further consideration. First, the current simulation did not involve the possibility of feedback from the adult agent (see, e.g. de Boer 2000). This factor would undoubtedly increase the learner's chances of attaining the target grammar. Morphological decomposability of lexical items would also contribute to the maintenance of the contrast in less favorable environments, reinforced by other kinds of lexical/grammatical generalizations, specifically, output-output correspondence relations between allomorphs (Benua 2000; see Kochetov 2002). Implementation of these factors is currently under way. Further, positing an a priori set of context-specific constraints made the simulation more manageable, but it would be more desirable for the learner to infer these from the input (cf., Hayes 1999). Further work should aim to rely on more complex interactions in agent populations over time and on a model of human (child and adult) speech production and perception that is more realistic and language-independent. Thus, one of the limitations of the current simulation is in its reliance on language-particular production of palatalized labials (Flemming, p.c.; see Kochetov 2002 for a discussion). This problem can be addressed through the use of both synthesized speech and data from a number of languages. Also, the current simulation does not predict the commonly attested phenomenon of contrast enhancement (Flemming 1995). It is possible that this effect may arise from continuous multiple-agent interactions (cf., de Boer 2000) with meaningful lexical items used in more natural communication environments. Finally, a more realistic simulation should give more attention to higher-level lexical and grammatical generalizations (see Wedel, this volume) and should use a wider range of lexical items and phonological contrasts.

5. Conclusion

In sum, the goal of the simulated learning situation was to provide us with insight into the nature and the mechanism of the emergence of positional markedness asymmetries in palatalization, and in general. An important result of the simulation is that it is unnecessary to attribute an innate status to positional markedness scales, or perhaps even to markedness scales in general. These phonological structures can arise and evolve through relatively simple and repetitive speaker-listener interactions during the course of language acquisition.

References

Beckman J. 1997. Positional faithfulness, positional neutralization and Shona vowel harmony. *Phonology* 14: 1-46.

Benua, L. 2000. *Phonological relations between words.* New York, N.Y.: Garland Publishing.

Browman, C. P., and Goldstein, L. 1999. Competing constraints on intergestural coordination and self-organization of phonological structures. In *Les Cahiers de l'ICP, Bulletin de la Communication Parlée* 5. Grenoble: 25-34.

de Boer, B. 2000. Self-organization in vowel systems. *Journal of Phonetics* 28: 441-465.

Flemming, E. 1995. Auditory representations in phonology. PhD dissertation, UCLA.

Hale, M., and C. Reiss. 2000. Substance abuse and dysfunctionalism: Current trends in phonology. *Linguistic Inquiry* 31: 157-169.

Hayes, B. 1999. Phonetically-driven phonology: The role of Optimality Theory and inductive grounding. In M. Darnell, E. Moravscik, M. Noonan, F. Newmeyer, and K. Wheatly, eds., *Functionalism and Formalism in Linguistics*, Vol. I: General Papers, John Benjamins, Amsterdam: 243-285.

Hayes, B. 2001. to appear. Phonological acquisition in Optimality Theory: The early stages. In *Fixing priorities: Constraints in phonological acquisition*, eds. R. Kager, J. Pater, and W. Zonneveld. Cambridge: Cambridge University Press.

Hume, E., and K. Johnson. 2001. A model of the interplay of speech perception and phonology. In *The role of speech perception in phonology*, eds. E. Hume and K. Johnson. New York, N.Y.: Academic Press: 3-26.

Johnson, K. 2003. *Acoustic and auditory phonetics.* 2nd edition. Cambridge, MA: Blackwell.

Kauffman, S. 1995. *At home in the universe: The search for the laws of self-organization and complexity*. Oxford: Oxford University Press.

Kochetov, A. 2002. *Production, perception, and emergent phonotactic patterns: A case of contrastive palatalization*. New York, London: Routledge.

Kochetov, A. to appear. Syllable position effects and gestural organization: Articulatory evidence form Russian. In *Papers in Laboratory Phonology VIII,* Berlin, New York: Mouton de Gruyter.

Langton, C. 1995. *Artificial Life: An overview*. MIT Press, Cambridge, MA.

Liberman, M. to appear. Simple models for emergence of a shared vocabulary. In *Papers in Laboratory Phonology VIII,* Berlin, New York: Mouton de Gruyter.

Lombardi, L. 1995. Laryngeal neutralization and syllable well-formedness. *Natural Language and Linguistic Theory* 13: 39-74.

Mielke, J. this volume. The diachronic influence of perception: Experimental evidence from Turkish. *Proceedings of the 29th meeting of the Berkeley Linguistic Society* (BLS 29).

Ohala, J.J. 1981. The listener as a source of sound change. In *Papers from Chicago Linguistic Society parasession on language and behavior*, C.S. Masek, R.A. Hendrick, and M.F. Miller. Chicago, IL: CLS: 178-203.

Prince, A., and P. Smolensky 1993. Optimality Theory: Constraint interaction in generative grammar. Ms., Rutgers University.

Steriade, D. 1997. Phonetics in phonology: The case of laryngeal neutralization. Ms., University of California, Los Angeles.

Steriade, D. 2001. The phonology of perceptibility effects: The P-map and its consequences for constraint organization. Ms.

Tesar, B., and P. Smolensky. 1998. Learnability in Optimality Theory. *Linguistic Inquiry* 29: 229-268.

Tesar, B. 2000. *Learnability in Optimality Theory*. Cambridge, MA: MIT Press.

Wedel, A. this volume. Self-organization and the origin of higher-order phonological patterns. *Proceedings of the 29th meeting of the Berkeley Linguistic Society* (BLS 29).

Alexei Kochetov
Department of Linguistics
Simon Fraser University
Burnaby, British Columbia
V5A 1S6 Canada

alexei_kochetov@sfu.ca