Learning artificial phonology: A review

(Authors' names suppressed) [*]

Version 3.6, April 25, 2011

Contents

---

[*] (Acknowledgements suppressed.)

1 Introduction

Laboratory study of the acquisition of artificial phonological patterns is potentially very significant as a window on the inductive biases involved in acquiring natural-language phonology.[1] Some such biases must exist, since generalizations cannot be learned without them (e.g., Pinker, 1979; Mitchell, 1990; Gallistel et al., 1991); what is at issue is rather their content and their causal relationship to phenomena of natural language acquisition, change, and typology. Artificial phonology may offer a way to study these biases in near-isolation from each other and from other confounding factors. The present article reviews the empirical literature on artificial-phonology learning in the context of this program.

The scope of this review is the acquisition artificial analogues of categorical phonology, i.e., patterns which partition a discrete stimulus space into positive ("legal", "pattern-conforming") versus negative ("illegal", "non-pattern-conforming") instances on the basis of phonological features. In the broader psychological literature, such partitions are often referred to as "concepts". We are not concerned here with the partitioning of a continuous stimulus space into phonetic categories (e.g., Maye et al., 2002; Goudbeek et al., 2008), nor with analogues of lexical (e.g., Peña et al., 2002; Perruchet et al., 2004; Newport and Aslin, 2004) or syntactic (e.g., Gómez, 2002) dependencies. Our interpretations of the experiments which we do review are not necessarily those of the original authors.

What factors make phonological patterns harder to learn? And do these highly artificial tasks reveal anything about natural-language phonology? Two biasing factors have been studied the most intensively, *formal complexity* and *phonetic substance.* The results of the review corroborate the early conclusions of Pycha et al. (2003): There is ample evidence that complexity impedes learning, but the picture for substantive bias — roughly speaking, a learning advantage for specific relationships between specific phonological features — is unclear. The relevant studies are thinly spread, and their results are highly variable. Natural-language typology and productivity are consistent with the effects of both complexity and substantive bias.

2 Formal complexity

Many formal theories of natural-language phonology are designed to favor patterns which have a simple expression in phonetic terms, and some proposals expressly impute this bias to human learners (e.g., Chomsky and Halle 1968, 330–334; Kiparsky 1971, 623; Bach and

---

[1] By *inductive bias* (also called "analytic bias" or "learning bias"), we mean any tendency of a pattern-learning algorithm to acquire one pattern faster or better than another when both are instantiated equally well by the training input. This definition is deliberately broad. It includes absolute distinctions between learnable and utterly unlearnable patterns, as well as relative distinctions between easier and harder ones. It is indifferent to details of implementation, applying to explicit penalties against specific patterns as well as to emergent consequences of the learner's architecture, representational system, or similarity metric. The term as we use it includes, but is not limited to, anything that would qualify as "Universal Grammar".

Harms 1972; Sagey 1990, 1; Hayes 1999; Gordon 2004, 304). Two main formal complexity factors have been studied in phonology: the number of features relevant to the pattern, and the relations between them.

2.1 Number of relevant features

A stimulus space described by three binary-valued features can be divided into two equal-sized categories in only six ways, if we ignore trivial variants obtained by permuting features or inverting feature values. Examples are shown in (1) for stimuli that are geometric figures varying in color (black vs. white), shape (circle vs. triangle), and size (large vs. small). Only color matters for the Type I distinction. Type II requires attention to color and shape, but size can be ignored. Types III through VI involve all three features, but some subsets can be decided with fewer (e.g., white triangles). For Type VI, not even this is possible; even a subset requires all three features.

(1) *Representatives of the six possible equal partitions of a stimulus space defined by binary features of color, shape, and size. Boxes enclose the (arbitrary) positive class. Concepts are arranged in increasing order of difficulty, with III, IV, and V being about equal. (After Shepard et al. 1961.)*

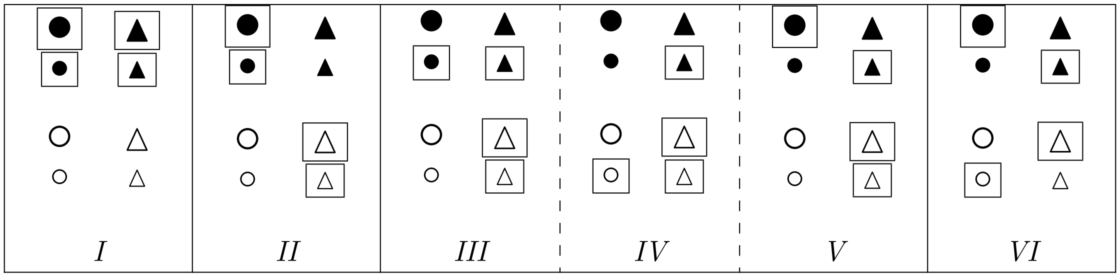

These six concepts have been extensively studied in connection with supervised learning of non-linguistic categories. In a typical experiment, the participant is shown a randomly-selected stimulus, judges whether it belongs to the target concept, and is then told the correct response. This cycle repeats until some performance criterion is met. The main finding is that difficulty increases along with the number of relevant features: Type I is easier than Type II, which is easier than Types III, IV, and V, which in turn are easier than Type VI (Shepard et al., 1961; Haygood and Bourne, 1965; Nosofsky et al., 1994; Feldman, 2000; Love, 2002).

Analogous phonological stimuli have been used in both supervised and unsupervised learning experiments. Participants are either trained with feedback to divide stimuli into “legal” and “illegal” categories, or are familiarized without feedback on “legal” stimuli only. They are then asked to categorize stimuli as legal or illegal, or to decide which of two stimuli is the legal one. The same difficulty hierarchy has been found for Types I, II, and VI. The other types have not to our knowledge been studied.

A phonological pattern that depends on a single stimulus feature (Type I) has often proven easier, and never harder, than one that requires more. Saffran and Thiessen (2003, Exps. 2, 3) familiarized English-learning 9-month-olds with isolated positive nonword instance, exposed them to a continuous stream of two positive and two negative nonwords, and then compared listening times to these four nonwords using headturn preference. When the pattern restricted [p t k] to some positions and [b d g] to others, the negative stimuli were preferred, but when the pattern distinguished [p d k] vs. [b t g], there was no difference in means (see schematic in Figure 2a). Cristiá and Seidl (2008, Exp. 1) familiarized English-learning 7-month-olds on positive C_1VC_2 nonwords. When the pattern was "C_1 is a nasal or oral stop" (i.e., [$-$continuant]), the infants preferred novel negative instances over positive ones; when it was "C_1 is a nasal or fricative", for which there is no standard feature, they showed no preference (Figure 2b). LaRiviere et al. (1974, 1977) trained English-L1 adults to categorize a set of six or eight syllables into two equal-sized classes defined either by a single feature or in an unspecified "random" way that needed more relevant features. Performance was significantly better for the single-feature condition than the random condition in three out of ten experiments, and was numerically better in the other seven.

(2) *Examples of phonological patterns in experiments comparing featural complexity. Some are defective representatives of their category types, owing to the impossibility of some feature combinations. In each case, the left-hand pattern proved the easier.*

(a) Saffran and Thiessen (2003, Exps. 2, 3). Features were voiced vs. voiceless, coronal vs. non-coronal, and labial vs. velar. (Defective.)

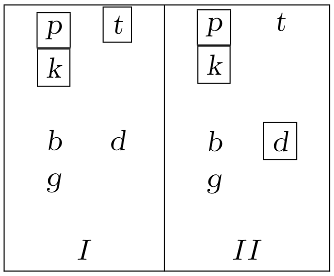

(b) Cristiá and Seidl (2008): Features were oral vs. nasal, continuant vs. non-continuant, labial vs. non-labial. (Defective.)

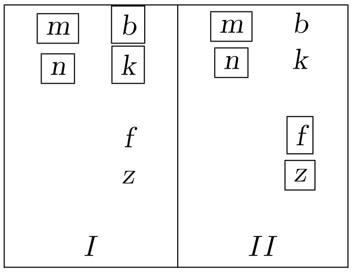

(c) *Kuo (2009): Features were plain initial stop vs. aspirated, labial initial stop vs. coronal, and palatal glide vs. labiovelar. Corresponding conditions (not shown) inverted the legal/illegal categories.*

$\boxed{pj}$	tj
$\boxed{p^hj}$	t^hj
pw	$\boxed{tw}$
p^hw	$\boxed{t^hw}$

II

pj	$\boxed{tj}$
$\boxed{p^hj}$	t^hj
$\boxed{pw}$	tw
p^hw	$\boxed{t^hw}$

VI

$\boxed{pj}$	$\boxed{tj}$
p^hj	t^hj
pw	tw
$\boxed{p^hw}$	$\boxed{t^hw}$

II

(d) *Pycha et al. (2003): Features are front first vowel vs. back, front last vowel vs. back, and high-lax first vowel vs. other first vowel.*

$\boxed{\text{i. . . ɛk}}$	u. . . ɛk
$\boxed{\text{æ. . . ɛk}}$	a. . . ɛk
$\boxed{\text{ɪ. . . ɛk}}$	ʊ. . . ɛk
i. . . ʌk	$\boxed{\text{u. . . ʌk}}$
æ. . . ʌk	$\boxed{\text{a. . . ʌk}}$
ɪ. . . ʌk	$\boxed{\text{ʊ. . . ʌk}}$

II

$\boxed{\text{i. . . ɛk}}$	u. . . ɛk
$\boxed{\text{æ. . . ɛk}}$	a. . . ɛk
ɪ. . . ɛk	$\boxed{\text{ʊ. . . ɛk}}$
i. . . ʌk	$\boxed{\text{u. . . ʌk}}$
æ. . . ʌk	$\boxed{\text{a. . . ʌk}}$
$\boxed{\text{ɪ. . . ʌk}}$	ʊ. . . ʌk

VI

i. . . ɛk	$\boxed{\text{u. . . ɛk}}$
æ. . . ɛk	$\boxed{\text{a. . . ɛk}}$
ɪ. . . ɛk	$\boxed{\text{ʊ. . . ɛk}}$
$\boxed{\text{i. . . ʌk}}$	u. . . ʌk
$\boxed{\text{æ. . . ʌk}}$	a. . . ʌk
$\boxed{\text{ɪ. . . ʌk}}$	ʊ. . . ʌk

II

Two-feature Type II patterns enjoy a similar advantage over three-feature Type VI ones. Kuo (2009) familiarized L1 Mandarin speakers on syllables with two-consonant onsets in which the initial stop perfectly predicted whether the following glide would be [j] or [w] (Figure 2c). In two patterns, a single stop feature, aspiration or place, was relevant (Type II); in the third, both stop features were needed (Type VI). Both Type II conditions elicited a significantly greater preference for novel positive stimuli over non-conforming foils than did the Type VI condition. A similar result was found by Pycha et al. (2003) in an experiment in which adult English speakers were trained with feedback to make binary grammaticality judgments of $X \ldots XY$ stimuli, where Y was either [-ɛk] or [-ʌk] depending on the vowel of X (Figure 2d). Classification at test was more accurate for two Type II patterns (backness agreement and backness disagreement) than for a Type VI pattern.

The advantage for patterns with fewer relevant features extends to patterns which are in part phonetically arbitrary. Using a speeded-repetition paradigm, Chambers et al. (2010) familiarized English speakers with a pattern in which the unsystematic sets [b f k m t] and [p s g n tʃ] were restricted to opposite ends of a C_1VC_2 syllable when the nucleus was one of two vowels, but unrestricted when it was a third vowel. In four experiments with different vowel sets, novel probe syllables were repeated faster when their consonants obeyed the positional restriction, regardless of what the vowel was. Participants evidently did not detect the dependency between V and the presence of positional restrictions on the Cs. Their performance followed an inaccurate moderately-complex generalization rather than an accurate more-complex one.

A complexity disadvantage has also been reported for learned alternations in production. Peperkamp et al. (2006) exposed French-speaking adults to stimuli of the form XY paired with pictures of two or three of the same object. The number of objects determined X ([nɛl] or [ʀa]), and the identity of the object determined Y. The initial consonant C_Y of each Y varied depending on X. In two conditions, a phonetically-defined set of consonants switched voicing ([f ʃ]↔[v ʒ], or [p k]↔[b g]). In two others, both the sets and the change were phonetically unsystematic ([p z]↔[ʒ f], or [ʃ v]↔[b k]). When tested on XY phrases with novel Ys, participants in the voicing conditions changed C_Y in the pattern-conforming way about 25% of the time, whereas those in the unsystematic conditions did so only about 5% of the time (most responses left C_Y unchanged). The relevant complexity here seems to be that of the change rather than that of the segment classes undergoing it: Participants in the voicing conditions did not generalize the rule to new segments in the old classes. They must have induced two single-segment rules rather than a class-based one (e.g., [f]↔[v] and [ʃ]↔[ʒ], rather than "voiceless fricatives alternate with voiced ones"). Evidently a rule like [f]↔[v] is learned better than one like [ʃ]↔[b]. Using similar stimuli with a similar population, Skoruppa et al. (2009) found that an alternation in which only place of articulation changed was learned better than one in which place and manner changed together, and also better than another in which place, manner, and voice changed together.

2.2 Relations between features

A separate question is whether certain syntagmatic relations between features *within* a stimulus facilitate pattern learing when the number of relevant features is controlled. Two main relations have been studied in artificial phonology, featural agreement and the contiguity-similarity tradeoff.

2.2.1 Featural agreement

Many phonological theories make special provision for representing agreement or disagreement of features within some part of an utterance (e.g., Chomsky and Halle, 1968; Goldsmith, 1976; Alderete and Frisch, 2008). Domain-general theories of category learning have not addressed this possibility, since they lack the means to recognize two features in the same stimulus as instances of the same abstract feature (Gluck and Bower, 1988; Kruschke, 1992; Nosofsky et al., 1994; Love et al., 2004; Feldman, 2006). However, there is non-linguistic evidence that within-stimulus featural identity is a relatively salient relation (Hunt and Hovland 1960; Ciborowski and Cole 1973; Ciborowski and Price-Williams 1974; not found by Shepard et al. 1961). The analogous evidence in phonological learning is suggestive, but not conclusive.

Wilson (2003) familiarized English speakers to stimuli of the form $C_1V_1C_2V_2C_3a$. The identity of C_2 determined whether C_3 was [n] or [l], as shown in (3).

(3) *Conditions of Experiments 1 and 2 of Wilson (2003).*

	Familiarization			*% judged familiar*			
	Nasal C_2	*Dorsal* C_2	*Other* C_2	*Old*		*Novel*	
Exp.	4×	4×	12×	*Conf.*	*Non.*	*Conf.*	*Non.*
1A	$C_1V_1[m/n]V_2na$	$C_1V_1[k/g]V_2la$	$C_1V_1C_2V_2la$	*70*	*44*	*53*	*34*
1B	$C_1V_1[m/n]V_2la$	$C_1V_1[k/g]V_2na$	$C_1V_1C_2V_2la$	*60*	*54*	*46*	*38*
2A	$C_1V_1[m/n]V_2la$	$C_1V_1[k/g]V_2na$	$C_1V_1C_2V_2na$	*73*	*47*	*50*	*35*
2B	$C_1V_1[m/n]V_2na$	$C_1V_1[k/g]V_2la$	$C_1V_1C_2V_2na$	*68*	*52*	*47*	*41*

The patterns in Conditions 1A and 2A can be stated as agreement or disagreement in [nasal] between C_2 and C_3, whereas those in 1B and 2B crucially involve a relation between two features [Dorsal] and [nasal], or [Dorsal] and [lateral]. Concept membership significantly increased judged familiarity in the single-feature Conditions 1A and 2A, but not in the two-feature Conditions 1B and 2B.[2] [3]

---

[2] Conditions 1A and 2A partially confound featural identity with segmental identity in the case of [n]; however, the difference between the C_2 = [m] and C_2 = [n] sub-conditions was not significant (Colin Wilson, p.c., 2010).

[3] This interpretation hinges on the traditional assumption that English post-tonic intervocalic [l] is phonologically [Coronal] despite its phonetic dorsal component (Sproat and Fujimura, 1993, 304). If [l] is phono-

Healy and Levitt (1980, Experiment 3) found that a voicing-conditioned pattern was acquired better by English speakers than a phonetically arbitrary one when the pattern was voicing assimilation, but not when it related voicing to vowel quality ([a] vs. [o]).[4] Similarly, Moreton (2008) familiarized English speakers on $C_1V_1C_2V_2$ stimuli and tested discrimination between novel positive and negative instances. Performance was better when the pattern was height agreement between the vowels, or voice agreement between the consonants, than when it was correlation between the height of V_1 and voicing of C_2.[5] Lin (2009), using the same stimuli with speakers of Mandarin and Southern Min, found better performance on the height-height than height-voicing pattern.

On the other hand, Kuo (2009) found no difference between a place-place correlation (labial glide iff labial stop) and a place-aspiration correlation (labial glide iff aspirated stop; see the two Type II patterns in Figure 2c). Seidl and Buckley (2005, Experiment 2) familiarized 9-month-old infants on $C_1V_1C_2V_2(C_3)$ stimuli and tested listening preference for novel positive vs. negative stimuli. A novelty preference was obtained for an agreement pattern in which C_1 and V_1 agreed in labiality, but also for one in which C_1 was labial if and only if V_1 was high.

2.2.2 Contiguity-similarity tradeoff

Phonological theory typically treats dependencies between adjacent elements as the normal case, excluding long-distance interactions unless the interacting segments share some property which is absent from intervening material (Jensen, 1974; McCarthy, 1981; Cole and Trigo, 1988; Pierrehumbert, 1993; Odden, 1995; Gafos, 1996; Hansson, 2001a; Frisch et al., 2004; Rose and Walker, 2004; Heinz, 2010). As with the other formal complexity biases reviewed above, there are parallels in non-linguistic learning: Two stimuli, or two elements of a compound stimulus, are more likely to cohere in perception and become associated in memory if they are contiguous in time or space, or are perceptually similar (Köhler 1941; Prentice and Asch 1958; Asch 1969; Arnold and Bower 1972; Rescorla 1980; Rescorla and Gillan 1980; Creel et al. 2004; Rescorla 2008; but see Pacton and Perruchet 2008.) It would therefore be surprising if contiguity and similarity did not facilitate acquisition of lab-learned phonology. However, the relevant evidence is scanty.[6]

Majerus et al. (2004) familiarized French speakers on a continuous stream of CV syllables which contained phonetically unsystematic CV and $C \ldots C$ dependencies, then tested

---

logically [Dorsal] as well (Walsh Dickey, 1997, Ch. 2), then all four conditions can be stated as single-feature agreement or disagreement, and this study may exemplify a substantive rather than a complexity bias.

[4] The "arbitrary" patterns were constructed in the same way as those of [ʃaffranθhiessenʌʉʉʒ], by replacing the classes [p t k] vs. [b d g] with [p t g] vs. [b d k].

[5] Featural identity was partially confounded with segmental identity, but the effect of featural identity was significant even for the subset of stimuli in which the segments were not identical.

[6] We omit here studies of artificial long-distance lexical (e.g. Peña et al., 2002; Perruchet et al., 2004; Newport and Aslin, 2004) and syntactic (e.g. Gómez, 2002) dependencies, as outside the scope of this review.

immediate recall of novel isolated nonword probes. Probes which belonged to both patterns simultaneously were recalled better than those which belonged to neither or only one, but there was no evidence that the two differed in effect. Using a tongue-twister paradigm, Warker and Dell (2006, 2008) tested English speakers on stimuli in which two consonants were constrained to appear at opposite ends of the stimulus. The positional restrictions were reversed depending on a third segment which was either adjacent to the marginal consonants ($C_1\mathbf{V}C_2$) or remote from them ($C_1V\mathbf{C}VC_2$). Exchange errors followed the positional restrictions to almost the same extent in the adjacent and remote conditions, with a numerical but nonsignificant advantage for the adjacent condition.

2.3 Summary: Complexity

There is abundant converging evidence that formal complexity impedes acquisition of artificial phonological patterns, in the sense that performance drops as the number of relevant features increases. The strongest result is that Type I patterns are easier than Type II, which are easier than Type VI. There is also evidence that stimulus-response mappings which change fewer features are easier than those which change more, and that within-stimulus dependencies are easier when they involve two instances of the same feature (i.e., assimilation or dissimilation) than instances of two different features. If the same inductive biases affect natural-language phonology, they should leave visible marks on cross-linguistic typology and within-language productivity.

A complexity bias in natural-language phonology would make more complex patterns harder to learn, hence harder to innovate and more likely to be changed (simplified) in tradition (Bach and Harms, 1972). That in turn could lead to low long-term steady-state frequencies for the corresponding patterns (Bell, 1970, 1971; Greenberg, 1978). Phonologists have in fact noted informally that the patterns they discover tend to be featurally simple (e.g., Chomsky and Halle 1968, 401, Hayes 1999, Pierrehumbert 2001). The available quantitative evidence tends to confirm this observation.

Clements (2003) found that inventories tend to avoid both “holes” and “bumps”: A given segment is more likely if all of its feature values are shared by other segments, and less likely if some of them are not. This is analogous to the Shepard et al. (1961) difficulty hierarchy, in which the difficult concepts of Types III–V involve isolated exceptions to simple featural rules (see Figure 4). In an iterated learning experiment simulating cultural transmission of non-linguistic Shepard-like concepts, Griffiths et al. (2008) have shown that such holes and bumps tend to be smoothed away over time.

(4) *The probability that a segment will occur in an inventory increases if the inventory contains other segments minimally different from it. (Extrapolated from Clements, 2003, Figure 11.)*

Favored				*Disfavored*	
$\boxed{p}$	$\boxed{t}$	$\boxed{p}$	$\boxed{t}$	$\boxed{p}$	$\boxed{t}$
$^*b$	$^*d$	$\boxed{b}$	$\boxed{d}$	$\boxed{b}$	$^*d$
$\boxed{f}$	$\boxed{s}$	$^*f$	$^*s$	$^*f$	$^*s$
$^*v$	$^*z$	$^*v$	$^*z$	$^*v$	$\boxed{z}$
Type I		*Type I*		*Type V*	

In a survey of 561 languages, (Mielke, 2004) studied "phonologically active classes", sets of sounds that pattern together by undergoing an alternation, triggering an alternation, or respecting a phonotactic restriction. One finding was that typologically common sound classes can usually be stated as a single feature value or a conjunction of a small number of feature values (.e.g, [–continuant –sonorant]), with typological frequency falling as the feature count rises. Of the non-conjunctive classes, most can be stated as disjunctions of conjunctions (e.g., [–sonorant Labial]∨[–sonorant Dorsal]); frequency falls as the number of disjuncts rises. Thus, featurally-complex patterns are attested but rarer, in the same way in which their artificial analogues are learnable but harder.[7]

An alternative hypothesis attributes the prevalence of simple classes to sampling error Pierrehumbert (2001, 2003). Suppose the learner decides whether to postulate a constraint by observing which of two classes, A or B, is more frequent. Classes defined by more features are rarer ($[+F+G]$ cannot outnumber $[+F]$), so if A and B both involve many features, the corpus of relevant examples will be small. For example, the learner can make a more reliable frequency comparison between $A = \{$ plosives $\}$ and $B = \{$ fricatives $\}$ than between $A = \{$ labial plosives $\}$ and $B = \{$ labial fricatives $\}$. Since there is more variability between smaller samples, learners will disagree more in their judgment of the relative frequency of A and B, and hence also in the constraints they acquire. That makes highly specific ("fine-grained") constraints are less likely to survive traditional transmission than very general ones. This idea may explain why conjunctive categories with more features are typologically rarer. However, it does not explain why *dis*junctive classes become rarer as the number of disjuncts goes up, since more disjuncts mean a larger class. E.g., the learner can make an even more reliable frequency comparison between $A = \{$ plosives and nasals $\}$ and $B = \{$ fricatives and laterals $\}$.

The simplicity of natural-language phonological patterns could also be inherited from their phonetic precursors: If interactions between phonetic variables in the speaker-hearer channel are usually simple, then phonological innovations inspired by those interactions ("phonologizations", Hyman 1976; Ohala 1993) will also tend to be simple. This hypothe-

---

[7] Many of the natural-language classes in the Mielke (2004) study involve more relevant features than the artificial-phonology experiments. However, they agree where they overlap. A Type I problem uses a single-feature class. A Type II problem requires a disjunction of two two-feature conjunctions, e.g., "(black and triangle) or (not-black and non-triangle)", while Type VI needs four three-feature disjuncts (Feldman, 2000, 2006; Lafond et al., 2007).

sis is not yet testable because there is no Mielke (2004)-like quantitative data on the typical complexity of phonetic precursors. We can ask whether inheritance is the *only* relevant factor by comparing the typological frequencies of patterns whose precursors have equal magnitude but different featural complexity. Where this has been tried, the more-frequent pattern was found to be associated with the simpler precursor, indicating that the typological skew towards simplicity was at least partly due to inductive bias (Moreton 2008, Moreton 2010; for an alternative suggestion see Yu 2010).

3 Phonetic substance

Patterns of equal formal complexity can differ widely in typological frequency. For example, coronal-stop assibilation is asymmetrically triggered by following rather than preceding vocoids, and by high rather than low vocoids (Kim, 2001); vowel-height harmony is more common than consonant continuancy harmony (Hansson 2001b, 137–149, Rose and Walker 2004). Very often, the more-frequent phonological pattern is a stylized version of some kind of phonetic covariation (e.g., phonological vowel harmony resembles phonetic vowel-to-vowel coarticulation). One family of explanations proposes inductive bias favoring patterns which have phonetic motivations over those which do not (e.g., Stampe, 1973; Prince and Smolensky, 1993; Archangeli and Pulleyblank, 1994; Hayes, 1999; Steriade, 2001; Wilson, 2006). The artificial-phonology experiments typically begin compare a typologically-frequent, phonetically-motivated pattern to a rare, unmotivated (or even counter-motivated) pattern of the same complexity. The available studies can be divided into three categories: consonant/vowel asymmetries, segmental rules, and prosodic rules.

3.1 Consonants versus vowels

There have been several reports that a within-stimulus vowel-identity dependency is learned faster than the analogous consonant pattern. Toro et al. (2008) familiarized native Italian speakers on $CVCVCV$ stimuli. In one condition, the first and last vowel were identical; in the other, the first and last consonant. Participants in the vowel condition preferred novel positive stimuli over non-conforming foils; those in the consonant condition did not (the two conditions were not directly compared to each other). The finding was replicated by Toro et al. (2008), where it was also shown to be robust against manipulation of the relative amplitudes of the consonants and vowels (again without direct comparison). It was replicated by Pons and Toro (2010) with Spanish-learning 16-month-olds; however, a direct comparison between the two conditions found no significant difference. Nevins and Toro (Nevins, 2010), in an experiment with Italian-speaking adults, directly compared consonant- with vowel-repetition patterns, and found a stronger preference for positive instances with the latter.

On the other hand, Koo (2007, Ch. 2), using a speeded-repetition paradigm with English speakers, found a conformity advantage for an [l. . . l]/[r. . . r] pattern and an [l. . . r]/[r. . . l] pattern, but not for analogous patterns with [i] and [u]. Since the responses in this experiment are utterances rather than judgments, the result may be due to the articulatory difficulty of co-occurring liquids, rather than to differences in learnability of the patterns. Two subsequent experiments in the same series with [l. . . l]/[r. . . r] and [i. . . i]/[u. . . u] found no difference in their effects on pattern-membership judgments of new stimuli.

On the whole, the experimental evidence is consistent with the hypothesis that learners are more sensitive to syntagmatic repetition of vowels than of consonants. If there is such an inductive bias, and if that bias shapes typology, we should find that patterns of non-adjacent vowel repetition outnumber the analogous consonant patterns in natural languages. Many languages have a phonological pattern in which one vowel is required to be identical to another, notably total vowel harmony (Aoki, 1968) and copy-vowel epenthesis (Kitto and de Lacy, 1999), whereas the consonantal analogues seem to us to be much rarer. However, we know of no quantitative test of this hypothesis.

3.2 Segmental rules

Wilson (2006) focused on two typological asymmetries in rules changing velars [k ɡ] to palatoalveolars [tʃ dʒ] as a function of vowel context. One is that palatalization before more-back vowels implies palatalization before less-back ones; the other, that palatalization of voiced velars implies that of voiceless ones. English speakers were trained in a language game to respond to a subset of [ki ɡi ke ɡe] with [tʃi dʒi tʃe dʒe], and to both of [ka ɡa] with [ka ɡa] (the critical syllables occurred initially in disyllabic nonsense words). They were then tested on a mix of old and new stimuli to measure their velar-palatalization rate in different conditions. Experiment 1 focused on the effect of vowel context; Experiment 2, that of consonant voicing. A synopsis is given in Table (5).

(5) *Critical experimental conditions of Wilson (2006). n, number of stimuli; p, probability of velar palatalization (in the training stimuli or in the test responses).*

			ki		*ke*		*ka*		*gi*		*ge*		*ga*	
Exp	*Cond*	*Phase*	*n*	*p*	*n*	*p*	*n*	*p*	*n*	*p*	*n*	*p*	*n*	*p*
1	*i*	*Train*	*4*	*1.00*	–	–	*3*	*0.00*	*4*	*1.00*	–	–	*3*	*0.00*
		Test	*8*	*0.44*	*8*	*0.13*	*6*	*0.05*	*8*	*0.52*	*8*	*0.14*	*6*	*0.14*
1	*e*	*Train*	–	–	*4*	*1.00*	*3*	*0.00*	–	–	*4*	*1.00*	*3*	*0.00*
		Test	*8*	*0.20*	*8*	*0.19*	*6*	*0.15*	*8*	*0.48*	*8*	*0.49*	*6*	*0.39*
2	*k*	*Train*	*4*	*1.00*	*4*	*1.00*	*3*	*0.00*	*1*	*1.00*	*1*	*1.00*	*3*	*0.00*
		Test	*8*	*0.39*	*8*	*0.36*	*6*	*0.12*	*8*	*0.14*	*8*	*0.11*	*6*	*0.09*
2	*g*	*Train*	*1*	*1.00*	*1*	*1.00*	*3*	*0.00*	*4*	*1.00*	*4*	*1.00*	*3*	*0.00*
		Test	*8*	*0.26*	*8*	*0.20*	*6*	*0.00*	*8*	*0.50*	*8*	*0.44*	*6*	*0.23*

The rate of palatoalveolar responses to velar-initial stimuli was bimodal, with clusters around 15% (Low) and 45% (Medium), rates comparable to those found by Peperkamp et al. (2006). Changing the features made a clear difference in performance in Experiment 1. Participants trained to respond [tʃe dʒe ka ga] (1e) had the Medium rate of palatoalveolar responses to [gi ge ga], but the Low rate to [ki ke ka]. Participants trained on [tʃi dʒi ka ga] (1i) had the Medium rate only on [ki] and [gi] themselves. These results are to some extent consistent with typology, since velar palatalization before [e] asymmetrically implies velar palatalization before [i].

However, there are also differences. Participants in Condition 1e, unlike natural-language palatalization rules, disregarded vowel context entirely. Their palatalization rates were indistinguishable before [e], where they had been trained to palatalize, [a], where they had been trained *not* to palatalize, and [i], where they had been given no training (Wilson, *op. cit.*, Fig. 2). In natural language, [g]-palatalization asymmetrically implies [k]-palatalization. If this is a result of inductive bias, we expect the learner to interpret observed [g]-palatalization as evidence for [k]-palatalization, but not the other way around. This was not borne out in the experiments. In Experiment 1, [g] was palatalized significantly more often than [k], despite equal training on both. In Experiment 2, there was no significant difference between the effects of [k]-training on [g]-palatalization and that of [g]-training on [k]-palatalization.

Vowel harmony is cross-linguistically very common, and a few studies have looked for special sensitivity to it. As described above (§2.1), Pycha et al. (2003) found no difference in learnability between artificial analogues of backness harmony and backness disharmony, whereas Moreton (2008) found better performance with a height-harmony analogue and a voice-harmony analogue than with a control condition (§2.2.1 above). Since backness disharmony and voice harmony are typologically rarer than backness harmony and height harmony, these findings amount to null results for an inductive bias for vowel harmony.

In typology, rounding harmony in mid vowels asymmetrically implies rounding harmony in high vowels (Kaun, 2004). In an experimental analogue, Finley and Badecker (2009, Exp. 3) familiarized English speakers, by passive listening, with stimuli of the form $X \ldots XY$ (where X was a $CVCV$ nonsense word and Y a CV suffix), and then asked them to choose the positive member of a pair of new stimuli XZ_1, XZ_2. The vowel of Y agreed in backness and rounding with those of X. For one group of participants, the vowels of Y were high and those of the Z's were mid; the reverse was true for the other group. Although familiarized participants chose the positive stimuli much more often than did unfamiliarized control listeners, the rate did not differ between the two familiarization groups. In other words, no analogue of the asymmetrical implication in natural language typology was found.

Zaba (2008, Ch. 2) used a paradigm similar to that of Pycha et al. (2003) to compare artificial analogues of three patterns differing in natural-language typological frequency: backness agreement between non-adjacent vowels (common), nasality agreement between non-adjacent consonants (rare), and labiality ageement between non-adjacent consonants (unattested in adult language). English-speaking participants were familiarized by listening to stimuli of the form $X \ldots XY$, where X ended in $\ldots V_1C_1$, and Y could be either of two V_2C_2 sequences. The pattern deterimined the choice of Y as a function of V_1 (for backness agreement) or C_1 (for nasality and labiality agreement). Learning was then tested by asking participants to judge whether $X \dots XY$ sequences conformed to the trained pattern. A control group of participants was likewise tested after being familiarized with $X \dots XY$ stimuli in which Y was always the same. The results were analyzed in several different ways, none of which found significant differences between the three pattern conditions.

3.3 Prosodic rules

Carpenter (2005, 2006, 2010) investigated acquisition of artificial stress patterns by native speakers of American English and Quebec French. Participants were trained to choose between correctly- and incorrectly-stressed versions of the same word, then tested with novel words. One set of experiments used sonority-sensitive stress, comparing the typologically uncommon “leftmost low vowel ([æ a]), else leftmost vowel” with the completely unattested “leftmost high vowel ([i u]), else leftmost vowel”. Participants in both native-language groups preferred positive items more strongly in the leftmost-low condition. Here the lab results are aligned with typology.[8]

A second set of experiments used quantity-sensitive stress, comparing the typologically frequent “leftmost heavy (CVC), else leftmost” pattern with the unattested “leftmost light (CV), else leftmost”. Since the typology is much more strongly skewed here than in the case of sonority-sensitve stress, one would expect the same of learning performance. However, although preference for positive stimuli was significantly above chance in all conditions, it did not differ significantly between the two artificial patterns in either native-language group.

Schane et al. (1974) trained participants to translate English adjective-plus-noun phrases word for word into an artificial language which had a context-sensitive rule deleting the final consonant of the first word (Figure 6). In one condition, deletion applied when the second word began with a consonant, simulating a cross-linguistically common process of cluster simplification (Wilson, 2001). Deletion in the other condition occurred when the second word was vowel-initial. Such intervocalic deletion in nature rarely applies to the voiceless obstruents used in this experiment (Picard, 2003).

(6) *Artificial-language conditions used by Schane et al. (1974).*

---

[8] Although the vowels were equalized for duration and peak intensity to remove phonetic stress cues, the author cautions us that participants may have restored low vowels’ inherently greater duration and intensity by subvocalizing the stimuli (Carpenter, 2010). However, the unedited high vowels were tense, while the unedited low ones were lax. Editing lengthened high vowels and shortened low ones (Carpenter, 2006, 72), exaggerating the tense/lax contrast in a way that would encourage English speakers to stress the high vowels, not the low ones.

Condition	$/C_1\#C_2/$ /ˈtupak ˈsipu/	$/C_1\#V_2/$ /ˈtupak ˈoga/
Cluster simplification	$\rightarrow [\#C_2]$ [ˈtupa ˈsipu]	[ˈtupak ˈoga]
Intervocalic C deletion	[ˈtupak ˈsipu]	$\rightarrow [\#V_2]$ [ˈtupa ˈoga]

Participants were trained to a performance criterion. The cluster-simplification groups reached criterion before the intervocalic-deletion ones, and were less likely to erroneously give responses appropriate to the other group.

In both conditions, the presence of a consonant at the end of Word 1 was correlated with that at the beginning of Word 2. Why would the negative correlation be easier to learn than the positive one? One possibility is substantive bias against marked phonological structures such as hiatus or syllable codas. Alternatively, participants in the cluster-simplification condition could have learned the pattern "exactly one intervocalic consonant", while those in the intervocalic-deletion condition would have had to learn the disjunction "zero or two intervocalic consonants". The results could then be accounted for by a complexity bias, the relative difficulty of disjunctive categories (Bruner et al., 1956; Conant and Trabasso, 1964; Ciborowski and Cole, 1972).

3.4 Summary: Substance

Research on phonetic substance has been spread thinly across a wide range of phenomena, with few attempts at replication. The best-supported finding, replicated with different stimuli and with speakers of different ages and languages, that it is easier to induce the pattern "contains the same vowel twice" than "contains the same consonant twice" in $CVCVCV$ stimuli. For other segmentally- and prosodically-defined patterns, results have been mixed at best, even within the same study.

The relation to typology is especially puzzling. When complexity bias mismatches typological frequency, the disparity may be due to a conflict with channel bias. This explanation is not available when substantive bias mismatches typology, since channel bias is predicted to *reinforce* substantive bias. Yet the studies reviewed here mismatch typology as often as they match it, even though the biases tested in the experiments were inferred from typology.

4 General discussion

The overall picture that emerges is that the difficulty of learning a phonological pattern in the lab increases as a function of the number of relevant features, but is much less dependent on the phonetic content of those features. However, these conclusions cannot be accepted until several further issues are settled.

4.1 An analytic artifact?

Calculations of both complexity and substance are affected by the analyst's choice of featural primitives, and feature theories are deliberately engineered to encode substantive bias as complexity bias (McCarthy, 1988; Clements and Hume, 1995). This theoretical prejudice could inflate the apparent frequency of complexity biases at the expense of substantive ones. For example, when Cristiá and Seidl (2008) found that [m n b k] vs. [f z] was easier to learn than [m n f z] vs. [b k], we interpreted it as evidence that Type I patterns are easier than Type II (see 2b above). That interpretation depended on the absence of a feature distinguishing [m n f z] from [b k]. If the learner in fact has such a feature (e.g., [continuous airflow] or [stable spectrum]), then reluctance to use it would constitute a substantive bias against a Type I pattern.

Other experiments, however, resist such reanalysis. To convert the Saffran and Thiessen (2003) experiment from Type I vs. Type II into Type I vs. Type I, we would have to credit the learner with a feature that separates [p d k] from [b t g]. To convert Kuo (2009)'s Type II vs. Type VI into Type II vs. Type II would require a feature that separates [p^{h} t] from [p t^{h}], and Pycha et al. (2003)'s II/VI comparison would need one that separated [i æ ʊ] from [u a ɪ]. Such phonetically arbitrary, *post hoc* features cannot be universal; they could only have been learned in the experiment itself. In that case, the unifying principle behind the experimental results would be that a problem is harder when a relevant feature is more complex phonetically — a complexity bias, not a substantive one.

Moreover, the substantive-bias hypothesis is not just about classes of sounds, but about their behavior; not just about obstruents, but about their devoicing in syllable-final position. Substantive biases are hypothesized to be predictable, either forwards from the phonetic factors which cause them, or backwards from the typological asymmetries which they cause. Although the Type I patterns used by LaRiviere et al. (1974, 1977); Saffran and Thiessen (2003), and Cristiá and Seidl (2008) were built on typologically common sound classes, the patterns themselves were phonetically unmotivated and typologically rare (e.g., "onsets must be nasal or oral stops"). Substance therefore does not explain why they were learned better than equally unmotivated Type II patterns. Pycha et al. (2003) found no statistically significant difference in difficulty between a phonetically motivated, typologically common Type II pattern of vowel backness agreement and a phonetically less-motivated, typologically rare Type II pattern of backness disagreement, though both proved easier than a Type VI pattern. The three Type II patterns compared by Zaba (2008, Ch. 2) differed widely in phonetic motivation and typological frequency, but not in learning outcome.

4.2 Is artificial phonology phonology at all?

Perhaps participants are treating the artificial-phonology task as if it were a non-linguistic concept-learning task about red triangles or fictitious diseases. Use of the same domain-general cognitive processes would predict shared complexity biases and lack of substantive bias. It may be that such domain-general mechanisms are involved in natural-language phonology as well (Hume and Johnson, 2001). If so, then artificial phonology is informative about natural-language phonology after all. However, it is also possible that natural-language phonology is learned using a separate set of dedicated processes, and hence that artificial phonology is irrelevant to it. Several strands of evidence are germane to this hypothesis.

4.2.1 Shared complexity biases

The sharing of complexity biases between artificial phonology and non-linguistic category learning does not prove that they share a common mechanism to the exclusion of natural-language phonology. The currently known shared biases are few, and most are so generic that even radically different learning algorithms can share them (Gluck and Bower, 1988; Kruschke, 1992; Nosofsky et al., 1994; Love et al., 2004; Feldman, 2006). Nature sometimes assigns the same function to two circuits, though behavioral similarity and theoretical parsimony would lead us to expect one (e.g., Simons-Wiedenmaier et al., 2006).

Furthermore, there is typological evidence that natural-language phonology shares at least one complexity bias with artificial phonology. Phonological patterns relating the height of vowels in adjacent syllables outnumber those relating the height of a vowel to the voicing of a following consonant. The typological-frequency difference is not matched by a difference in size or frequency of the phonetic precursors, but artificial analogues of the two patterns differ in learnability in favor of the featural-agreement rule (§2.2.1) (Moreton, 2008; Lin, 2009). This is so far the only known case of its type, and so may be an anomaly.

4.2.2 Age and robustness of acquisition

The separate-mechanism hypothesis would be strengthened if there were an age at which infants can rapidly learn artificial phonology, but show no knowledge of native-language phonology. Infants start acquiring first-language phonotactic patterns between 6 and 9 months Jusczyk et al. (1993); Friederici and Wessels (1993); Jusczyk et al. (1994); Mattys and Jusczyk (2001), and can rapidly learn artificial phonotactics at the ages of 7–9 months (Marcus et al., 1999; Saffran and Thiessen, 2003; Seidl and Buckley, 2005; Cristiá and Seidl, 2008). The crucial experiments — artificial phonotactics in infants younger than 6 months — have not to our knowledge been done.

Most of the studies reviewed above used less than 30 minutes of familiarization or training, and some used much less. Taylor and Houghton (2005, Exp. 5) used a tongue-twister paradigm to familiarize English speakers with a pattern that restricted different consonants to different syllabic positions (onset vs. coda). They then reversed the pattern unannounced in the middle of a block. Before the switch, transposition errors adhered to the original constraints 98% of the time. After it, the new constraints were adhered to 65% of the time, with the violations concentrated in the first 9 trials after the switch. In contrast, first-language phonology is notoriously persistent, even in highly motivated speakers with ample exposure to second-language input (Cutler et al., 1989; Darcy, 2006; Kager et al., 2008). Artificial phonology thus appears easier to lose than first-language phonology and easier to acquire than second-language phonology, supporting the separate-mechanism hypothesis.

On closer inspection, the evidence becomes less clear. *Some* first-language patterns resist retraining, but others can be overcome with a small amount of lab exposure. Carpenter (2005, 2006, 2010) succeeded in teaching English speakers the artificial stress rule "leftmost low vowel, else leftmost vowel". Since the low vowels were also lax and the high ones also tense, the artificial pattern overcame the English L1 preference for stressing heavy syllables. Kuo (2009) found that Mandarin speakers' L1 phonotactics did not affect their ability to acquire L1-conforming vs. L1-nonconforming artificial patterns from brief exposure.

4.2.3 Interactions between natural and artificial phonology

Interactions between natural and artificial phonology are evidence that the two are not wholly separate processes. Schane et al. (1974)'s cluster-simplification pattern was, by design, similar to French liaison, and was acquired better by those who had had more exposure to French. Healy and Levitt (1980, Experiment 2) found better performance with a pattern that resembled English voicing agreement (in *-ed*, *-'s*, etc.) than for one that did not. Pater and Tessier (2006), again with English speakers, found better performance on an artificial epenthesis rule when it was triggered by violation of the English minimal-word constraint than when it was triggered by vowel backness. It is not clear whether natural language phonology had these effects by facilitating rule learning in training, or by biasing participants' responses against non-conforming responses in testing, but it is clear that some aspect of the artificial-phonology task are ecologically valid enough to engage natural-language phonological preferences.

5 Conclusions

The lab results so far add up to substantial and, in our view, convincing evidence that formal complexity impedes artificial-phonology learning in adults and infants, but the effect of phonetic substance is more elusive. These findings corroborate the suggestion of Pycha et al. (2003) that complexity bias is stronger than substantive bias in artificial phonology. If natural-language phonology works the same way, then the broader picture is essentially that proposed by Bach and Harms (1972): Inductive bias facilitates the faithful acquisition of simple patterns and the simplification of more complex ones, but is relatively insensitive to their phonetic motivation. This inductive bias acts as a filter on sources of new phonological patterns, such as ambient phonetic variation (Hyman, 1976; Ohala, 1992, 1993), children's spontaneous innovations (Stampe, 1973), and contact between languages or dialects (Trudgill, 1986). The outcome of historical change is thus jointly skewed by a combination of the source and filter biases. The resulting asymmetries in the direction of historical change are what determine long-term steady-state typological frequencies (Bell, 1970, 1971; Greenberg, 1978). Our companion paper (reference suppressed) discusses the computational modelling of simplicity bias in learning and its effect on typology, in both the phonological literature and in the psychological concept learning literature, which as we discussed in §2.1 above, has extensively examined analogous biases in that domain.

References

Alderete, J. and S. A. Frisch (2008). Dissimilation in grammar and the lexicon. In P. de Lacy (Ed.), *The Cambridge handbook of phonology*, Chapter 16, pp. 379–398. Cambridge, England: Cambridge University Press.

Aoki, H. (1968). Towards a typology of vowel harmony. *International Journal of American Linguistics 34*(2), 142–145.

Archangeli, D. and D. Pulleyblank (1994). *Grounded Phonology.* Cambridge, Massachusetts: MIT Press.

Arnold, P. G. and G. H. Bower (1972). Perceptual conditions affecting ease of association. *Journal of Experimental Psychology 93*(1), 176–180.

Asch, S. E. (1969). A reformulation of the problem of associations. *American Psychologist 24*(2), 92–101.

Bach, E. and R. T. Harms (1972). How do languages get crazy rules? In R. P. Stockwell and R. K. S. Macaulay (Eds.), *Linguistic change and generative theory*, Chapter 1, pp. 1–21. Bloomington: Indiana University Press.

Bell, A. (1970). *A state-process approach to syllabicity and syllabic structure.* Ph. D. thesis, Stanford University.

Bell, A. (1971). Some patterns of the occurrence and formation of syllabic structure. *Working Papers on Language Universals 6*, 23–138.

Bruner, J. S., J. J. Goodnow, and G. A. Austin (1956). *A study of thinking.* New York: John Wiley and Sons.

Carpenter, A. C. (2005). Acquisition of a natural vs. an unnatural stress system. In A. Brugos, M. R. Clark-Cotton, and S. Han (Eds.), *Papers from the 29th Boston University Conference on Language Development (BUCLD 29)*, Somerville, pp. 134–143. Cascadilla Press.

Carpenter, A. C. (2006). *Acquisition of a natural versus an unnatural stress system.* Ph. D. thesis, University of Massachusetts, Amherst.

Carpenter, A. C. (2010). A naturalness bias in learning stress. *Phonology 27*(3), 345–392.

Chambers, K. E., K. H. Onishi, and C. Fisher (2010). A vowel is a vowel: generalizing newly learned phonotactic constraints to new contexts. *Journal of Experimental Psychology:*

Learning, Memory, and Cognition 36(3), 821–828.

Chomsky, N. and M. A. Halle (1968). *The sound pattern of English.* Cambridge, Massachusetts: MIT Press.

Ciborowski, T. and M. Cole (1972). A cross-cultural study of conjunctive and disjunctive concept learning. *Child Development 43*, 774–789.

Ciborowski, T. and M. Cole (1973). A developmental and cross-cultural study of the influences of rule structure and problem composition on the learning of conceptual classifications. *Journal of Experimental Child Psychology 15*(2), 193–215.

Ciborowski, T. and D. Price-Williams (1974). The influence of rule structure and problem composition on conceptual learning among rural Hawaiian children. Technical Report 75, Kamehameha Early Education Program, Kamehamea Schools and Bernice P. Bishop Estate, 1850 Makuakane Street, Honolulu, Hawaii 96817.

Clements, G. N. (2003). Feature economy in sound systems. *Phonology 20*(3), 287–333.

Clements, G. N. and E. V. Hume (1995). The internal organization of speech sounds. In J. A. Goldsmith (Ed.), *The handbook of phonological theory*, Chapter 7, pp. 245–306. Boston: Blackwell.

Cole, J. and L. Trigo (1988). Parasitic harmony. In H. van der Hulst and N. Smith (Eds.), *Features, segmental structure, and harmony processes*, pp. 19–38. Foris.

Conant, M. B. and T. Trabasso (1964). Conjunctive and disjunctive concept formation under equal-information conditions. *Journal of Experimental Psychology 67*(3), 250–255.

Creel, S. C., E. L. Newport, and R. N. Aslin (2004). Distant melodies: statistical learning of non-adjacent dependencies in tone sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition 30*, 1119–1130.

Cristiá, A. and A. Seidl (2008). Is infants' learning of sound patterns constrained by phonological features? *Language Learning and Development 4*(3), 203–227.

Cutler, A., J. Mehler, D. Norris, and J. Segui (1989). Limits on bilingualism. *Nature 340*, 229–230.

Darcy, I. (2006). *Assimilation phonologique et reconnaissance des mots.* Publications Universitaires Européennes, Série XXI: Linguistique. Bern: Peter Lang.

Feldman, J. (2000). Minimization of Boolean complexity in human concept learning. *Nature 407*, 630–633.

Feldman, J. (2006). An algebra of human concept learning. *Journal of mathematical psychology 50*, 339–368.

Finley, S. and W. Badecker (2009). Artificial language learning and feature-based generalization. *Journal of Memory and Language 61*, 423–437.

Friederici, A. D. and J. M. I. Wessels (1993). Phonotactic knowledge of word boundaries and its use in infant speech perception. *Perception and Psychophysics 54*, 287–295.

Frisch, S., J. B. Pierrehumbert, and M. B. Broe (2004). Similarity avoidance and the OCP. *Natural Language and Linguistic Theory 22*(1), 179–228.

Gafos, A. (1996). *The articulatory basis of locality in phonology.* Ph. D. thesis, Johns

Hopkins University, Baltimore.

Gallistel, C. R., A. L. Brown, S. Carey, R. Gelman, and F. C. Keil (1991). Lessons from animal learning for the study of cognitive development. In S. G. Carey and R. G. Gelman (Eds.), *The epigenesis of mind*, Chapter 1, pp. 3–36. Hillsdale, New Jersey: Lawrence Erlbaum Associates.

Gluck, M. A. and G. H. Bower (1988). Evaluating an adaptive network model of human learning. *Journal of Memory and Language 27*, 166–195.

Goldsmith, J. A. (1976). *Autosegmental phonology.* Ph. D. thesis, Massachusetts Institute of Technology.

Gómez, R. (2002). Variability and detection of invariant structure. *Psychological Science 13*(5), 431–436.

Gordon, M. (2004). Syllable weight. In B. Hayes, R. Kirchner, and D. Steriade (Eds.), *Phonetically-based phonology*, pp. 277–312. Cambridge, England: Cambridge University Press.

Goudbeek, M., A. Cutler, and R. Smits (2008). Supervised and unsupervised learning of multidimensionally varying non-native speech categories. *Speech communication 50*, 109–125.

Greenberg, J. H. (1978). Diachrony, synchrony, and language universals. In J. H. Greenberg, C. A. Ferguson, and E. A. Moravcsik (Eds.), *Universals of human language, volume 1, method and theory*, pp. 61–91. Stanford, California: Stanford University Press.

Griffiths, T. L., M. L. Kalish, and S. Lewandowsky (2008). Theoretical and empirical evidence for the impact of inductive biases on cultural evolution. *Philosophical Transactions of the Royal Society B (Biological Sciences) 363*(1509), 3503–3514.

Hansson, G. Ó. (2001a). The phonologization of production constraints: Evidence from consonant harmony. In M. Andronis, C. Ball, H. Elston, and S. Neuvel (Eds.), *Papers from the 37th meeting of the Chicago Linguistics Society (CLS 37)*, Volume 1, Chicago, pp. 187–200. Chicago Linguistic Society.

Hansson, G. Ó. (2001b). *Theoretical and typological issues in consonant harmony.* Ph. D. thesis, University of California, Berkeley.

Hayes, B. (1999). Phonetically driven phonology: the role of optimality in inductive grounding. In M. Darnell, E. Moravcsik, M. Noonan, F. Newmeyer, and K. Wheatly (Eds.), *Functionalism and Formalism in Linguistics*, Volume 1: General Papers, pp. 243–285. Amsterdam: John Benjamins.

Haygood, R. C. and L. E. Bourne (1965). Attribute- and rule-learning aspects of conceptual behavior. *Psychological Review 72*(3), 175–195.

Healy, A. F. and A. G. Levitt (1980). Accessibility of the voicing distinction for learning phonological rules. *Memory and Cognition 8*(2), 107–114.

Heinz, J. (2010). Learning long-distance phonotactics. *Linguistic Inquiry 41*(4), xx–xx.

Hume, E. and K. Johnson (2001). A model of the interplay of speech perception and phonology. In E. Hume and K. Johnson (Eds.), *The role of speech perception in phonology*, pp. 3–26. San Diego: Academic Press.

Hunt, E. B. and C. L. Hovland (1960). Order of consideration of different types of concepts. *Journal of Experimental Psychology 59*(4), 220–225.

Hyman, L. M. (1976). Phonologization. In A. Juilland (Ed.), *Linguistic studies offered to Joseph Greenberg: second volume: phonology*, pp. 407–418. Saratoga, California: Anma Libri.

Jensen, J. T. (1974). A constraint on variables in phonology. *Language 50*(4, Part 1), 675–686.

Jusczyk, P. W., A. D. Friederici, J. M. I. Wessels, V. Y. Svenkerud, and A. M. Jusczyk (1993). Infants' sensitivity to the sound pattern of native language words. *Journal of Memory and Language 32*, 402–420.

Jusczyk, P. W., P. A. Luce, and J. Charles-Luce (1994). Infants sensitivity to phonotactic patterns in the native language. *Journal of Memory and Language 33*, 630–645.

Kager, R., N. Boll-Avetisyan, and C. Ao (2008, November). Gradient phonotactic constraints for speech segmentation in a second language. Poster presented at the Boston University Conference on Language Development.

Kaun, A. (2004). The typology of rounding harmony. In B. Hayes, R. Kirchner, and D. Steriade (Eds.), *Phonetically based phonology*, Chapter 4, pp. 87–116. Cambridge, England: Cambridge University Press.

Kim, H. (2001). A phonetically based account of phonological stop assibilation. *Phonology 18*(1), 81–108.

Kiparsky, P. (1971). Historical linguistics. In W. Dingwall (Ed.), *A survey of linguistic science*, pp. 576–653. College Park: University of Maryland.

Kitto, C. and P. de Lacy (1999). A Correspondence theory of epenthetic quality. In C. Kitto and C. Smallwood (Eds.), *Proceedings of AFLA (Austronesian Formal Linguistics Association VI*, Toronto Working Papers in Linguistics, pp. 181–200. Toronto: Unspecified publisher.

Köhler, W. (1941). On the nature of associations. *Proceedings of the American Philosophical Society 84*, 489–502.

Koo, H. (2007). *Change in the adult phonological processing system by learning non-adjacent phonotactic constraints from brief experience: an experimental and computational study.* Ph. D. thesis, University of Illinois, Urbana-Champaign.

Kruschke, J. K. (1992). ALCOVE: an exemplar-based connectionist model of category learning. *Psychological Review 99*, 22–44.

Kuo, L. (2009). The role of natural class features in the acquisition of phonotactic regularities. *Journal of psycholinguistic research 38*(2), 129–150.

Lafond, D., Y. Lacouture, and G. Mineau (2007). Complexity minimization in rule-based category learning: revising the catalog of Boolean concepts and evidence for non-minimal rules. *Journal of Mathematical Psychology 51*, 57–75.

LaRiviere, C., H. Winitz, J. Reeds, and E. Herriman (1974). The conceptual reality of selected distinctive features. *Journal of Speech and Hearing Research 17*(1), 122–133.

LaRiviere, C., H. Winitz, J. Reeds, and E. Herriman (1977). Erratum: The conceptual reality of selected distinctive features. *Journal of Speech and Hearing Research 20*(4), 817.

Lin, Y. (2009). Tests of analytic bias in native Mandarin speakers and native Southern Min speakers. In Y. Xiao (Ed.), *21st North American Conference on Chinese Linguistics*, Smithfield, Rhode Island, pp. 81–92. Bryant University.

Love, B. C. (2002). Comparing supervised and unsupervised category learning. *Psychonomic Bulletin and Review 9*(4), 829–835.

Love, B. C., D. L. Medin, and T. M. Gureckis (2004). SUSTAIN: a network model of category learning. *Psychological Review 111*(2), 309–332.

Majerus, S., L. Mulder, T. Meulmans, and F. Peters (2004). Verbal short-term memory reflects the sublexical organization of the phonological language network: evidence from an incidental phonotactic learning paradigm. *Journal of Memory and Language 51*, 297–306.

Marcus, G. F., S. Vijayan, S. B. Rao, and P. M. Vishton (1999). Rule learning by seven-month-old infants. *Science 283*, 77–80.

Mattys, S. L. and P. W. Jusczyk (2001). Phonotactic cues for segmentation of fleunt speech by infants. *Cognition 78*(2), 91–121.

Maye, J., J. F. Werker, and L. Gerken (2002). Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition 82*, B101–B111.

McCarthy, J. J. (1981). A prosodic theory of nonconcatenative morphology. *Linguistic Inquiry 12*, 373–418.

McCarthy, J. J. (1988). Feature Geometry and dependency: a review. *Phonetica 45*, 84–108.

Mielke, J. (2004). *The emergence of distinctive features.* Ph. D. thesis, Ohio State University.

Mitchell, T. M. (1980 [1990]). The need for biases in learning generalizations. In J. W. Shavlik and T. G. Detterich (Eds.), *Readings in machine learning*, Chapter 2.4.1, pp. 184–191. San Mateo, California: Morgan Kaufman.

Moreton, E. (2008). Analytic bias and phonological typology. *Phonology 25*(1), 83–127.

Moreton, E. (2010). Underphonologization and modularity bias. In S. Parker (Ed.), *Phonological argumentation: essays on evidence and motivation*, Chapter 3, pp. 79–101. London: Equinox.

Nevins, A. (2010). Two case studies in phonological universals: a view from artificial grammars. *Biolinguistics 4*(2), 218–233.

Newport, E. and R. N. Aslin (2004). Learning at a distance i: statistical learning of non-adjacent dependencies. *Cognitive Psychology 48*, 127–162.

Nosofsky, R. M., M. A. Gluck, T. J. Palmeri, S. C. McKinley, and P. Gauthier (1994). Comparing models of rule-based classification learning: a replication and extension of Shepard, Hovland, and Jenkins (1961). *Memory and Cognition 22*(3), 352–369.

Nosofsky, R. M., T. J. Palmeri, and S. C. McKinley (1994). Rule-plus-exception model of classification learning. *Psychological Review 101*(1), 53–79.

Odden, D. (1995). Adjacency parameters in phonology. *Language 70*(2), 289–330.

Ohala, J. J. (1992). What's cognitive, what's not, in sound change. In G. Kellerman and M. D. Morissey (Eds.), *Diachrony within synchrony: language history and cognition*, pp. 309–355. Frankfurt am Main: Peter Lang.

Ohala, J. J. (1993). The phonetics of sound change. In C. Jones (Ed.), *Historical linguistics: problems and perspectives*, pp. 237–278. Harlow: Longman.

Pacton, S. and P. Perruchet (2008). An attention-based associative account of adjacent and nonadjacent dependency learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition 34*(1), 80–96.

Pater, J. and E. Moreton (2011). Simplicity bias in phonological learning. Manuscript, University of Massachusetts at Amherst and University of North Carolina at Chapel Hill.

Pater, J. and A.-M. Tessier (2006). L1 phonotactic knowledge and the L2 acquisition of alternations. In R. Slabakova, S. Motrul, P. Prévost, and L. White (Eds.), *Inquiries in linguistic development: in honor of Lydia White*, pp. 115–131. Benjamins.

Peña, M., L. L. Bonatti, M. Nespor, and J. Mehler (2002). Signal-driven computations in speech processing. *Science 298*, 604–607.

Peperkamp, S., K. Skoruppa, and E. Dupoux (2006). The role of phonetic naturalness in phonological rule acquisition. In D. Bamman, T. Magnitskaia, and C. Zoller (Eds.), *Papers from the 30th Boston University Conference on Language Development (BUCLD 30)*, pp. 464–475.

Perruchet, P., M. D. Tyler, N. Galland, and R. Peereman (2004). Learning nonadjacent dependencies: no need for algebraic-like computations. *Journal of Experimental Psychology: General 133*(4), 573–583.

Picard, M. (2003). On the emergence and resolution of hiatus. *Folia Linguistica Historica 24*(1–2), 47–57.

Pierrehumbert, J. B. (1993). Dissimilarity in the Arabic verbal roots. In A. Schafer (Ed.), *Proceedings of the Northeast Linguistics Society*, Volume 23, Amherst, Mass., pp. 367–381. Graduate Linguistics Students Association.

Pierrehumbert, J. B. (2001). Why phonological constraints are so coarse-grained. *Language and Cognitive Processes 16*(5–6), 691–698.

Pierrehumbert, J. B. (2003). Probabilistic phonology: discrimination and robustness. In R. Bod, J. Hay, and S. Jannedy (Eds.), *Probabilistic phonology*, Chapter 6, pp. 177–228. Cambridge, Massachusetts: MIT Press.

Pinker, S. (1979). Formal models of language learning. *Cognition 7*, 217–283.

Pons, F. and J. M. Toro (2010). Structural generalizations over consonants and vowels in 11-month-old infants. *Cognition 116*, 361–367.

Prentice, W. C. H. and S. E. Asch (1958). Paired associations with related and unrelated pairs of nonsense-figures. *American Journal of Psychology 71*, 247–254.

Prince, A. and P. Smolensky (1993). *Optimality Theory: constraint interaction in generative grammar.* Department of Linguistics, Rutgers University.

Pycha, A., P. Nowak, E. Shin, and R. Shosted (2003). Phonological rule-learing and its implications for a theory of vowel harmony. In M. Tsujimura and G. Garding (Eds.), *Proceedings of the 22nd West Coast Conference on Formal Linguistics (WCCFL 22)*, pp. 101–114.

Rescorla, R. A. (1980). Simultaneous and successive associations in sensory preconditioning. *Journal of Experimental Psychology: Animal Behavior Processes 6*(3), 207–216.

Rescorla, R. A. (2008). Evaluating conditioning of related and unrelated stimuli using a compound test. *Learning and Behavior 36*(2), 67–74.

Rescorla, R. A. and D. J. Gillan (1980). An analysis of the facilitative effect of similarity on second-order conditioning. *Journal of Experimental Psychology: Animal Behavior Processes 6*(4), 339–351.

Rose, S. and R. Walker (2004). A typology of consonant agreement as correspondence. *Language 80*(3), 475–531.

Saffran, J. R. and E. D. Thiessen (2003). Pattern induction by infant language learners. *Developmental Psychology 39*(3), 484–494.

Sagey, E. (1990). *The representation of features in non-linear phonology: the Articulator Node Hierarchy.* New York: Garland.

Schane, S. A., B. Tranel, and H. Lane (1974). On the psychological reality of a natural rule of syllable structure. *Cognition 3*(4), 351–358.

Seidl, A. and E. Buckley (2005). On the learning of arbitrary phonological rules. *Language Learning and Development 1*(3 & 4), 289–316.

Shepard, R. N., C. L. Hovland, and H. M. Jenkins (1961). Learning and memorization of classifications. *Psychological Monographs 75*(13, Whole No. 517).

Simons-Wiedenmaier, N. S., M. Weber, C. F. Plappert, P. K. D. Pilz, and S. Schmid (2006). Synaptic depression and short-term habituation are located in the sensory part of the mammalian startle pathway. *BMC Neuroscience 7*(38).

Skoruppa, K., A. Lambrechts, and S. Peperkamp (2009). The role of phonetic distance in the acquisition of phonological alternations. In S. Lima, K. Mullin, and B. Smith (Eds.), *Proceedings of the 39th Meeting of the North-East Linguistics Society (NELS)*, Amherst, Massachusetts. Graduate Linguistics Students Association.

Sproat, R. and O. Fujimura (1993). Allophonic variation in English /l/ and its implications for phonetic implementation. *Journal of Phonetics 21*, 291–311.

Stampe, D. (1973). *A dissertation on Natural Phonology.* Ph. D. thesis, University of Chicago.

Steriade, D. (2001). The phonology of perceptibility effects: the P-map and its consequences for constraint organization. MS, Department of Linguistics, University of California, Los Angeles.

Taylor, C. F. and G. Houghton (2005). Learning artificial phonotactic constraints: time course, durability, and relationship to natural constraints. *Journal of Experimental Psychology: Learning, Memory, and Cognition 31*, 1398–1416.

Toro, J. M., M. Nespor, J. Mehler, and L. L. Bonatti (2008). Finding words and rules in a speech stream. *Psychological Science 19*(2), 137–144.

Toro, J. M., M. Shukla, M. Nespor, and A. D. Endress (2008). The quest for generalizations over consonants: asymmetries between consonants and vowels are not the by-product of acoustic differences. *Perception and Psychophysics 70*(8), 1515–1525.

Trudgill, P. (1986). *Dialects in contact.* Oxford, England: Blackwell.

Walsh Dickey, L. (1997). *The phonology of liquids.* Ph. D. thesis, University of Massachusetts, Amherst, Massachusetts.

Warker, J. A. and G. S. Dell (2006). Speech errors reflect newly learned phonotactic constraints. *Journal of Experimental Psychology: Learning, Memory, and Cognition 32*, 387–398.

Warker, J. A., G. S. Dell, C. A. Whalen, and S. Gereg (2008). Limits on learning phonotactic constraints from recent production experience. *Journal of Experimental Psychology: Learning, Memory, and Cognition 34*(5), 1209–1295.

Wilson, C. (2001). Consonant cluster neutralisation and targeted constraints. *Phonology 18*, 147–197.

Wilson, C. (2003). Experimental investigation of phonological naturalness. In G. Garding and M. Tsujimura (Eds.), *Proceedings of the 22nd West Coast Conference on Formal Linguistics (WCCFL 22)*, Somerville, pp. 533–546. Cascadilla Press.

Wilson, C. (2006). Learning phonology with substantive bias: an experimental and computational study of velar palatalization. *Cognitive Science 30*(5), 945–982.

Yu, A. C. L. (2010). Tonal effects on perceived vowel duration. In C. Fougeron, B. Kühnert, M. D'Imperio, and N. Vallée (Eds.), *Papers in Laboratory Phonology 10: phonology and phonetics*, pp. 151–168. Mouton De Gruyter.

Zaba, A. M. (2008). *Relative frequency of patterns and learnability: the case of phonological harmony.* Ph. D. thesis, University of Utah.