

The Role of Phonetic Distance in the Acquisition of Phonological Alternations*

Katrin Skoruppa¹, Anna Lambrechts¹ and Sharon Peperkamp¹²

¹Laboratoire de Sciences Cognitives et Psycholinguistique, EHESS-DEC(ENS)-CNRS, Paris

²Université de Paris 8

1. Introduction

In phonology, it has often been noted that some regularities are typologically more frequent than others. For example, alternations that change [i] to [u] occur more often than alternations that change [i] into [ɨ] (Chomsky and Halle, 1968). In order to explain such typological asymmetries, many phonological theories refer to their phonetic bases: Typologically frequent regularities tend to ‘make sense’ from a phonetic point of view (e.g. Donegan et al., 1979; Archangeli and Pulleybank, 1994). Two proposals as to how **phonetic naturalness** gives rise to typological asymmetries are the following. First, cognitive biases have been argued to play a role during phonological acquisition in each individual language learner; these biases typically favor the acquisition of phonetically natural regularities (e.g. Jusczyk, 1998; Wilson, 2006). Second, sound changes that take place over generations of speakers have been argued to yield more phonetically natural regularities than unnatural ones (e.g. Ohala, 1992; Blevins, 2004). Thus, phonetically natural regularities might be typologically frequent because they are easier to acquire (synchronic perspective) or because they arise more frequently (diachronic perspective).¹

Several recent studies have investigated experimentally whether phonetic naturalness has an effect on phonological learning. In these studies, participants are exposed to highly controlled artificial languages. Typically, one group of participants is exposed to stimuli that follow phonological patterns considered to be phonetically natural, whereas another group is exposed to unnatural patterns. The performances of both groups during a subsequent test phase are then compared in order to assess learning differences. It should be noted that although it has often been treated as a monolithic, one-dimensional concept, the term ‘phonetic naturalness’ comprises several aspects (Peperkamp et al., 2006b). So far, artificial language-learning studies have focused on two of them:

First, most phonological alternations involve **feature spreading**, or assimilation, such that the output of an alternation is phonetically closer to one or more neighbouring sounds than the input. An example is vowel harmony, a widespread phonological process that makes

* This work was funded by a grant from the Agence Nationale de la Recherche (ANR 05 BLAN-0065-01). We are grateful to Anne-Caroline Fiévet and Inga Vendelin for help in recruiting participants, and to Audrey Breton and Alicia Coste for help in running participants. We also thank Anne Christophe and the other members of the LSCP for discussion.

¹ These two proposals are not mutually exclusive; see Moreton (2008) for recent discussion.

all vowels of a word share a phonetic feature. In Hungarian, for instance, vowels in roots and suffixes must have the same place of articulation, thus causing suffix vowel alternations: The dative form of “*ház*” (‘house’) is “*ház-nak*” [ha:znak], but the dative form of “*öröm*” (‘joy’) is “*öröm-nek*” [øømnək]. Several studies have explored whether alternations that involve feature spreading are easier to learn than those that do not, with divergent results: Wilson (2003) showed that adult participants were better at learning a consonant harmony rule than an arbitrary rule on consonant co-occurrence. Neither Pycha et al. (2003), nor Skoruppa and Peperkamp (submitted), however, reported a significant difference between the acquisition of vowel harmony and vowel disharmony by English-speaking adults. Seidl and Buckley (2005) found no effect of feature spreading in nine-month-old infants either; assimilatory and arbitrary constraints on consonant-vowel sequencing were indeed learned to the same extent.

A second aspect of phonetic naturalness concerns **natural classes**: Alternations typically apply to groups of sounds that share one or more phonetic features. For example, Hungarian vowel harmony changes *all* front vowels into corresponding back vowels and vice versa. Regarding the impact of natural classes on phonological learning, several studies found that alternations that apply to a single natural class of sounds are easier to learn than alternations that apply to arbitrary sound groupings. Saffran and Thiessen (2003), for instance, found that infants learned phonotactic patterns better if they applied to a natural class of sounds (here, voiceless stops), than if they did not. Similar results were found with a larger natural class (all stops) in another infant study (Cristià and Seidl, 2008). Finally, the studies by Pycha et al. (2003) and Skoruppa and Peperkamp (submitted) described above - in which adults successfully learned vowel harmony and vowel disharmony - both included a third group of participants who were exposed to an arbitrary combination of harmony and disharmony; crucially, this group did not learn the regularity in either study.

Another set of artificial language learning studies investigated the generalization of newly learned alternations to sounds that were not used during exposure, but belonged to the same natural class as the exposure sounds. Peperkamp et al. (2006b) and Peperkamp and Dupoux (2007) found that participants learn a new rule of intervocalic voicing of stops or fricatives, but do not generalize it to novel sounds of the same class. Wilson (2006), by contrast, found that participants not only learned a new rule of palatalization of [k] before [e] but also generalized the application of this process to [k] before [i].

In this article, we examine a third aspect of phonetic naturalness, that to our knowledge has not yet been investigated experimentally. As noted already by Trubetzkoy, (1939/1958:45) for allophonic variation, alternating sounds tend to be phonetically close to each other, that is, the change between them typically concerns a small number of features. This aspect of phonetic naturalness, **phonetic distance**, explains why phonetically dissimilar sounds that are in complementary distributions “by accident” are not allophones. For example, [h] and [ŋ] have complementary distributions in English because of syllable structure constraints, but they are too phonetically dissimilar to be considered as allophones.

Trubetzkoy (1939:32) observed that allophones are phonetically minimally distant; in particular, there is no other sound that shares all common features of the allophones. In the English example cited above, the glottal fricative consonant [h] and the velar nasal consonant [ŋ] cannot be considered as allophones, because the velar fricative consonant [x] shares all common features of [ŋ] and [h] (that is, being a consonant). Non-allophonic alternations tend to be phonetically minimal, too. Note that this minimal phonetic distance principle is discrete,

not gradient: there should be no sound between the alternating sounds, but the exact amount of phonetic distance between them (or the number of intervening sounds) does not matter.

As mentioned before, no study has yet addressed specifically the role of phonetic distance in phonological learning.² The present experiment examines whether phonetic distance has an effect on phonological learning, and if so, if the effect is discrete or gradient. To this end, we compare the learning of alternations that involve one, two, or three feature changes. If phonetic distance plays a role, these three types of alternations should yield different results. In particular, if phonetic distance has a discrete, all-or-nothing effect, one-feature changes should be learned better than two- and three-feature changes, whereas the latter two should yield the same results. If, by contrast, the effect of phonetic distance is gradient, one-feature changes should be learned better than two-feature changes, which in turn should be learned better than three-feature changes. Following Schane et al. (1974), we analyze the participants' responses not only during the test phase but also during the preceding learning phase. We thus address the question of whether phonetically minimal alternations are not only learned better, but also faster.

2. Experiment

Six artificial languages were constructed, sharing the same segmental inventory. This inventory is a subset of the one of the participants' native language, i.e. French. Each language has two obstruent alternations that do not exist in French (see Table 1).

Table 1: Sound alternations in the six languages used in the present experiment.

<i>Phonetic distance</i>	<i>Language</i>	<i>Alternating sounds</i>	
		<i>pair 1</i>	<i>pair 2</i>
Small (place)	S1	p - t	z - ʒ
	S2	ʃ - s	d - b
Medium (place and manner)	M1	p - s	d - ʒ
	M2	ʃ - t	z - b
Large (place, manner, and voicing)	L1	p - z	t - ʒ
	L2	ʃ - d	s - b

In the first two languages, the phonetic distance between the alternating consonants is small, that is, they differ in a single feature (place of articulation, e.g. [p]–[t]); two more languages contain medium-distance alternations involving two feature changes (place and manner, e.g. [p]–[s]); the last two languages contain large-distance alternations involving three feature changes (place, manner and voicing, e.g. [p]–[z]). Note that each of the consonants [p, b, d, t, s, z, ʃ, ʒ] participates in one small-distance, one medium-distance and one large-distance alternation.

The alternations in languages S1 and S2 satisfy the minimal phonetic distance principle, that is, there is no sound sharing all common features of the alternating sounds. By contrast, the alternations in the other four languages violate this principle, since for each alternation there is a consonant in the language that shares all common features of the two

² Peperkamp et al. (2006b) found that a one-feature change in consonant voicing only (e.g. [f]–[v]) was easier to learn than a three-feature change (e.g. [f]–[g]), but this parameter was not manipulated independently from the two other aspects that were examined, feature spreading and natural classes, making the origin of the naturalness effect hard to determine.

alternating ones (for example in M1, the alveolar stop [t] lies between the alternating sounds, the bilabial stop [p] and the alveolar fricative [s]). Thus, if phonetic distance has an impact on the learnability of alternations, languages S1 and S2 should be easier to learn than the others. If this effect is gradient, languages M1 and M2 should also be easier to learn than L1 and L2, whereas the latter four should be learned equally poorly if it is discrete.

2.1 Method

2.1.1 Stimuli

For each language, twelve pairs of phrases with alternating sounds were constructed. For each of these experimental phrase pairs, one phrase consisted of the monosyllabic non-word [ʔe] followed by a disyllabic non-word starting with one of the obstruents [p, b, d, t, s, z, ʃ, ʒ] (e.g. *ʔe pamu*); the other phrase consisted of a different monosyllabic non-word, i.e. [nø], followed by the same disyllabic non-word whose initial consonant, though, was changed (e.g. *nø tamu*). The change in consonant was determined by the language-specific alternations mentioned above. Examples of phrase pairs for the two alternations in each language are shown in Table 2. Note that except for the first consonant of the disyllables, each phrase pair is strictly identical across the six languages.

Table 2: Two examples of experimental phrase pairs in the six languages.

<i>Language</i>	<i>Alternating phrases</i>	
	<i>pair 1</i>	<i>pair 2</i>
S1	ʔe p amu – nø t amu	ʔe z fam – nø ʒ fam
S2	ʔe ʃ amu – nø s amu	ʔe d fam – nø b fam
M1	ʔe p amu – nø s amu	ʔe d fam – nø ʒ fam
M2	ʔe ʃ amu – nø t amu	ʔe z fam – nø b fam
L1	ʔe p amu – nø z amu	ʔe t fam – nø ʒ fam
L2	ʔe ʃ amu – nø d amu	ʔe s fam – nø b fam

Each alternation was present in six experimental phrase pairs for every language; a complete list is shown in Table 4 in the Appendix.³

Twelve more phrase pairs were constructed as above, except that the disyllabic non-word started with one of the sonorants [l,ʔ,m,n]. These filler phrase pairs were the same for all languages and did not contain alternations (e.g. *ʔe nibut* – *nø nibut*).

³ All alternations apply in a quite unnatural context (that is, the roundedness of the preceding vowel, [e] or [ø], triggers the alternation); we thus avoid a possible confound due to a feature spreading effect. The direction of change (fronting or backing) was counterbalanced across languages: In all languages of type 1, [ʔe] was combined with the more anterior consonant (e.g. *ʔe pamu* for S1) and [nø] with the more posterior consonant (e.g. *nø tamu* for S1) of an alternating pair. The opposite was true for type 2 languages. Finally, in order to avoid ease of articulation of certain obstruent-vowel combinations as a possible confound, the same phrases – modulo the voicing of the alternating consonants – were used for every given phonetic distance size. For instance, as shown in Table 2, *ʔe pamu* occurs in languages S1, M1, and L1, whereas *nø tamu* occurs in S1 and M2 and its voicing counterpart *nø damu* in L2. Likewise, *ʔe famu* occurs in S2, M2 and L2, whereas *nø samu* occurs in languages S2 and M1 and its voicing counterpart *nø zamu* in L1. Taken together, the S-languages thus contained the same phrases as the M-languages and the L-languages (with the proviso concerning voicing mentioned above).

All phrases were recorded naturally in random order by a phonetically trained native speaker of French. She pronounced the phrases with penultimate stress, such that they sounded more foreign to the French participants, who are used to final stress.

2.1.2 Procedure

Participants were tested on a computer in a quiet room wearing headphones. They were told that they would be hearing short phrases of the type ‘small dog’ or ‘big ship’ in an unknown language, in which *re* means ‘small’ and *nø* means ‘big’, and that their task would be to say aloud the same noun preceded by the opposite adjective. They were informed that sounds could change according to the adjectives.⁴

During pre-training, participants were familiarized with the task on six filler phrase pairs in the presence of an experimenter who gave further explanation, if necessary, and coded the correctness of their responses. Each trial was structured as follows:

Step 1: After a blank screen of one second, the word *Moi* : (‘I:’) appeared on the screen. One second later, the first phrase of a pair was played (e.g. *re nibut*).

Step 2: One second later, the word *Vous* : (‘You:’) appeared on a new line below the former, and the participant could give a response. She had to press a button when finished.

Step 3: The word *Correct* : (‘Correct:’) appeared on a new line, and one second later the second phrase of the pair was played (e.g. *nø nibut*).

Step 4: The experimenter coded via the mouse whether the participant’s response was correct. If so, she received positive feedback visually and moved on to the next trial; if not, she was told to try again and the trial was repeated.

In half of the trials, participants heard a phrase starting with [ʁe] (and had to produce the corresponding phrase starting with [nø]), in the other half, they heard a phrase starting with [nø]. At the end of the pre-training phase, the experimenter left the room, and the participant’s responses were henceforth recorded via microphone on a digital audio recorder for later transcription. Participants were informed that they would continue the same task and that they would still hear the correct response after their own one, but not given corrective feedback anymore.

The subsequent training phase contained a list of four experimental and two filler phrase pairs that was presented six times in different random orders, for a total of 36 trials. The trial procedure was the same as during pre-training, except for the on-line coding and the corrective feedback (hence, Step 4 was no longer part of the procedure). The filler pairs were the same for all participants, whereas the realization of the experimental pairs differed according to the language of exposure. For each language, two experimental pairs contained the first alternation (e.g. [p-t] for S1), and two the second one (e.g. [z-ʒ] for S1). For half of the trials, participants heard a phrase starting with [ʁe] (and had to produce the corresponding phrase starting with [nø]), in the other half, they heard a phrase starting with [nø]. For each pair, they had to produce both phrases three times.

⁴ In pilot studies in which participants were not informed about possible changes, they reported that they had considered the changes as recording errors and never applied any alternation themselves.

Finally, 36 test trials were presented. Participants were informed that they would not hear the correct answer anymore. A list of the six training items and twelve novel items (eight experimental items and four fillers) was played twice in random order with the same trial procedure as during training, except that the correct answer was no longer played (hence, Step 3 was no longer part of the procedure). For each language, half of the experimental items contained the first alternation and half the second one. For half of the items in each condition, the phrase started with [nø], for the others, the phrase started with [ɛe].

The experiment lasted about ten minutes.

2.1.3 Participants

Thirty-six young adults, all monolingual native French speakers without known history of language or hearing disorders, participated in the experiment. They were randomly assigned to one of six groups, with each of the groups being exposed to a different language.

2.2 Results and discussion

The participants' productions were transcribed by a native French speaker who was unaware of the phrases that had been presented. Responses with one of the following characteristics were discarded: incorrect monosyllable ([ɛe] or [nø]), unintelligible or missing disyllable (total: 13.5%). Stress errors were ignored. It was checked that participants had pressed the button to listen to the correct answer *after* they had given their response.

Figure 1 (overleaf) shows the progression of each individual participant during training for the experimental items. Here, the total number of correct responses that a given subject accumulated is plotted against the trial number, connected by a filled line. The diagonal dashed line depicts the best possible performance.

These learning curves are steeper for the participants exposed to languages with small-distance alternations than for those exposed to the other languages. An ANOVA with the between-subject factor 'phonetic distance' (small vs. medium vs. large) revealed a significant effect of this factor on the integrals, that is, the size of the surface below the individual curves ($F(2,33)=10.2$, $p<.001$). Planned pairwise comparisons by independent t-tests showed that the integrals of the small-distance group (122.8) are significantly larger than the ones of the medium-distance group (45.5, $t(20)=3.12$, $p<.01$) and the ones of the large-distance group (32.2, $t(16)=4.1$, $p<.001$). There was no significant difference between the latter two ($t(20)<1$). These analyses show that the alternations with small phonetic distance were learned faster than the two other types.

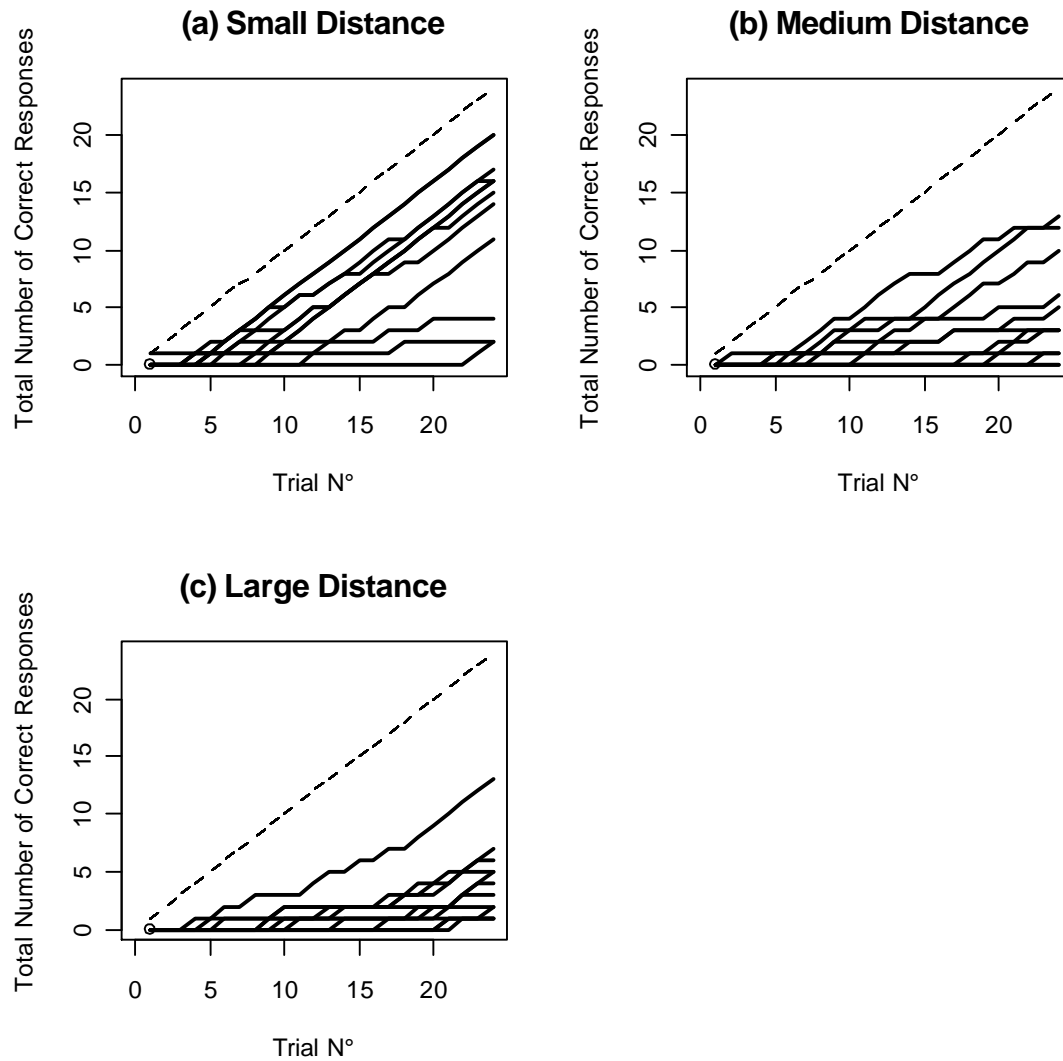


Figure 1: *Total number of correct responses by trial number for experimental items during training for individual participants exposed to alternations with small (a), medium (b), or large (c) phonetic distance*

Figure 2 (overleaf) shows the percentage of correct responses during the test phase for the experimental items by item type and by size of phonetic distance of the alternations.

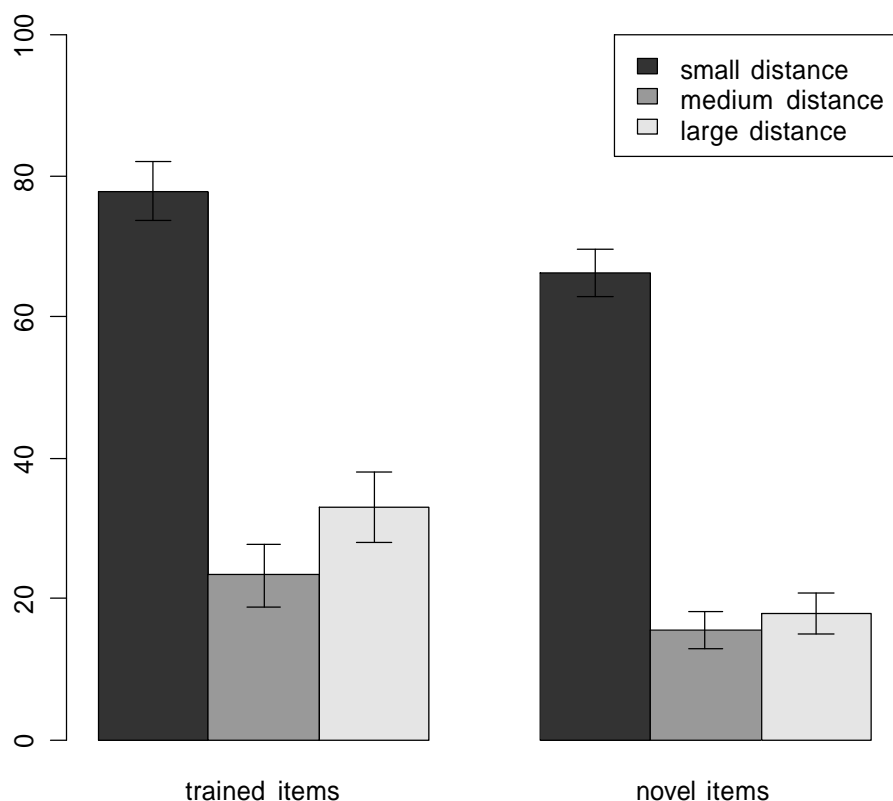


Figure 2: Mean percentages of correct responses for experimental items during test phase by phonetic distance size (left panel: trained items, right panel: novel items)

An ANOVA with the between-subject factor ‘phonetic distance’ (small vs. medium vs. large) and the within-subject factor ‘item type’ (trained vs. novel) revealed significant main effects of phonetic distance ($F(2,30)=10.3$, $p<.001$) and of item type ($F(2,33)=17.9$, $p<.001$), but no interaction ($F<1$). Planned pairwise comparisons of the three groups by independent t-tests were carried out separately for trained and novel items. For trained items, participants in the small-distance group gave significantly more correct responses (77.9%) than those in the medium-distance group (23.3%, $t(182)=8.8$, $p<.001$) and in the large-distance group (33.0%, $t(174)=6.8$, $p<.001$), whereas the latter two groups did not differ from one another ($t(173)=1.4$, $p=.16$). Similar effects were found for the novel items: Participants in the small phonetic distance group gave significantly more correct responses (66.3%) than those in the medium (15.4%, $t(354)=11.7$, $p<.001$) and the large phonetic distance groups (17.9%, $t(361)=10.9$, $p<.001$), whereas the latter two groups did not differ ($t(368)<1$) from one another.

Thus, alternations characterized by a small phonetic distance between the alternating sounds are learned faster and better than alternations characterized by a medium or a large phonetic distance between the alternating sounds. These results suggest that phonetic distance between alternating sounds influences their ease of acquisition. Moreover, the effect is not gradient, but discrete, with medium- and large-distance alternations with at least one intermediate segment being equally hard to learn.

Finally, we analyzed the error patterns of the different participant groups. Most erroneous responses consisted of the participant not changing the critical consonant at all (No Change errors, e.g. *ʁe pamu* in reply to *nø pamu*), as can be seen in Table 3.

Table 3: Mean percentages of No Change errors by phonetic distance size and item type for experimental items during test phase.

Item Type	Small	Medium	Large
Trained	76.2%	81.9%	66.1%
Novel	87.7%	82.8%	74.5%

An ANOVA with the between-subject factor ‘phonetic distance’ (small vs. medium vs. large) and the within-subject factor ‘item type’ (trained vs. novel) revealed no significant effects or interaction. The remaining errors were changes to other consonants, but their number was too small to carry out further analyses.

3. Conclusion

Using a rather explicit artificial language-learning paradigm, we found that all else being equal, phonetic distance affects the ease of acquisition of phonological alternations by adult learners. Its effect is discrete rather than gradient, in that phonetically non-minimal changes are hard to learn compared to minimal changes, regardless of the precise amount of phonetic distance involved.

One caveat is in order, though. Whereas the small-distance alternations in the present experiment concerned place changes, both the medium-distance and the large-distance alternations contained manner changes, which are acoustically rather salient. We cannot exclude that medium- and large-distance changes were harder to learn solely because of the presence of these manner changes. However, in a pilot study using a similar methodology, we showed that the presence versus absence of an acoustically less salient feature, i.e. place, has an impact on learning too.⁵ To further investigate this issue, it would be interesting to use a manner change for the small-distance alternation and add place and voice for the medium- and large-distance alternations. Unfortunately, given that in French most manner changes come with a place change (indeed, its stops are bilabial, dental and velar, whereas its fricatives are labiodental, dental and postalveolar), this is impossible with the current design.

It would also be interesting to lengthen the training phase, to explore whether with more exposure the differences between the three types of alternation would eventually disappear. Likewise, the experimental design could be adapted such as to use a perception rather than a production task. This would allow us to investigate whether the phonetic distance effect is specific to production. Previous work indeed showed that with identical

⁵ In the pilot study, twelve native French adults were exposed to small distance changes (voicing, e.g. *nel pima – ʁa bima*), and twelve others to medium distance changes (place and voicing, e.g. *nel pima – ʁa dima*) in a similar paradigm as the one used in the present study. Participants learning small distance changes gave more correct responses than participants learning medium distance changes, for trained items (57.3% vs. 39.8%, $t(262)=2.9$, $p<.01$) as well as for novel items (55.6% vs. 39.0%, $t(626)=4.2$, $p<.001$). These results, though, are possibly confounded by the fact that the context of the alternations is natural for voicing changes (which occur intervocalically) but unnatural for place changes, as well as by the fact that French has a voicing alternation (regressive voice assimilation in obstruents clusters), which might make French participants more sensitive to voicing than to place alternations.

exposure, a naturalness effect can be found when a production task is used (Peperkamp et al., 2006b) but not when a forced-choice perception task is used (Peperkamp and Dupoux, 2007).

The evidence so far that phonetic distance has a discrete effect on learnability meshes well with a computational study by Peperkamp et al. (2006a) that simulated the acquisition of allophonic rules based on the tracking of complementary distributions. In this study, results were greatly improved when a discrete, phonetic distance filter *à la* Trubetzkoy (1939) was added, discarding pairs of segments as possible allophones if an intermediate segment existed. This filter (together with a second one that discards non-assimilatory rules) dramatically reduced the number of segment pairs that were erroneously identified as allophones, such as the English pair [h] - [ŋ].

Of course, experimental studies with infants and toddlers are necessary to ultimately determine whether phonetic distance plays a role during the acquisition of native phonological rules (allophonic or not), and if so, whether the effect is discrete. In that case, rule learning would contrast with word recognition, for which phonetic distance has been shown to yield a *gradient* effect in toddlers. In particular, White et al. (2008) presented 19-month olds with mispronounced labels of familiar objects (e.g. ‘gall’ instead of ‘ball’). The phonetic distance between the substituted sound and the correct one was varied: One-feature (‘gall’), two-feature (‘zall’) and three-feature changes (‘shawl’) were investigated. Crucially, phonetic distance had a linear, gradient effect on toddlers’ looking times to the picture of the familiar object (a ball): The more features were changed, the less they fixated the picture, showing that they have graded sensitivity to phonetic distance, at least in familiar words.

To conclude, we have focused on a heretofore unexplored aspect of phonetic naturalness, that is, phonetic distance. The results of our experiment show that alternations that involve phonetically minimal distant sounds are learned faster and better than other alternations. They thus add to a growing body of evidence that phonetic naturalness influences the learnability of phonological alternations by adults as well as by infants (Pycha et al., 2003; Saffran and Thiessen, 2003; Wilson, 2003, 2006; Seidl and Buckley, 2005; Peperkamp et al., 2006b; Cristià and Seidl, 2008; Skoruppa and Peperkamp, submitted).

References

- Archangeli, D., and Pulleybank, D. 1994. *Grounded Phonology*. Cambridge, MA: MIT Press.
- Blevins, J. 2004. *Evolutionary Phonology: The emergence of sound patterns*. Cambridge: Cambridge University Press.
- Chomsky, N.A., and Halle, M. 1968. *The Sound Pattern of English*. New York: Harper and Row.
- Cristià, A., and Seidl, A. 2008. Is infants’ learning of sound patterns constrained by phonological features? *Language Learning and Development* 4:203-227.
- Donegan, P., Stampe, D., and Dinnsen, D. 1979. The Study of Natural Phonology. In *Current Approaches to Phonological Theory*, 126 -173. Bloomington, Indiana: Indiana University Press.
- Jusczyk, P. W. 1998. Constraining the search for structure in the input. *Lingua* 106:197-218.
- Moreton, E.. 2008. Analytic Bias as a Factor in Phonological Typology. *Phonology* 25:83-127.
- Ohala, J. J. 1992. What's cognitive, what's not, in sound change. In *Diachrony within synchrony: Language history and cognition*. [Duisburger Arbeiten zur Sprach- und

The role of phonetic distance in the acquisition of phonological alternations

- Kulturwissenschaft 14]*, eds. G. Kellermann and M. D. Morrissey, 309-355. Frankfurt am Main: Peter Lang Verlag.
- Peperkamp, S., Le Calvez, R., Nadal, J.-P., and Dupoux, E. 2006a. The acquisition of allophonic rules: statistical learning with linguistic constraints. *Cognition* 101: B31-B41.
- Peperkamp, S., Skoruppa, K., and Dupoux, E.. 2006b. The Role of Phonetic Naturalness in Phonological Acquisition. In *Proceedings of the 30th Annual Boston University Conference on Language Development*, eds. D. Bamman, T. Magnitskaia and C. Zaller, 464-475. Somerville: Cascadilla Press.
- Peperkamp, S., and Dupoux, E.. 2007. The acquisition of abstract phoneme categories. In *Laboratory Phonology 9*, eds. J. Cole and J. Hualde, 315-338. Berlin: Mouton de Gruyter.
- Pycha, A., Nowak, P., Shin, E., and Shosted, R. 2003. Phonological rule-learning and its implications for a theory of vowel harmony. In *Proceedings of WCCFL 22* eds. G. Garding and M. Tsujimura, 423-435. Somerville, Cascadilla Press.
- Saffran, J. R., and Thiessen, E. D. 2003. Pattern Induction by Infant Language Learners. *Developmental Psychology* 39:484-494.
- Schane, S. A., Tranel, B., and Lane, H. 1974. On the psychological reality of a natural rule of syllable structure. *Cognition* 3:351-358.
- Seidl, A., and Buckley, E.. 2005. On the learning of natural and unnatural rules. *Language Learning and Development* 1:289-316.
- Skoruppa, K., and Peperkamp, S. submitted. Learning a novel accent: Implicit acquisition of phonological alternations.
- Trubetzkoy, N. S.. 1939/1958. *Grundzüge der Phonologie*. Göttingen: Vandenhoeck & Ruprecht.
- White, K S., and Morgan, J. L. 2008. Sub-segmental detail in early lexical representations. *Journal of Memory and Language* 59:114-132
- Wilson, C. 2003. Experimental investigation of phonological naturalness. In *Proceedings of WCCFL 22* eds. G. Garding and M. Tsujimura, 101-114. Somerville: Cascadilla Press.
- Wilson, C. 2006. Learning phonology with substantive bias: An experimental and computational study of velar palatalization. *Cognitive Science* 30:945-982.

Laboratoire de Sciences Cognitives et Psycholinguistique (DEC-ENS)
29, rue d'Ulm
75005 Paris, France

skoruppa@ens.fr

Appendix: Experimental item pairs by language and item type

N°	Item Type	Small distance		Medium distance		Large distance	
		S1	S2	M1	M2	L1	L2
1	trained	ke p amu nø t amu	ke ʃ amu nø s amu	ke p amu nø s amu	ke ʃ amu nø t amu	ke p amu nø z amu	ke ʃ amu nø d amu
2	trained	ke z apiʃ nø ʒ apiʃ	ke d apiʃ nø b apiʃ	ke d apiʃ nø ʒ apiʃ	ke z apiʃ nø b apiʃ	ke t apiʃ nø ʒ apiʃ	ke s apiʃ nø b apiʃ
3	trained	ke z umi nø ʒ umi	ke d umi nø b umi	ke d umi nø ʒ umi	ke z umi nø b umi	ke t umi nø ʒ umi	ke s umi nø b umi
4	trained	ke p inaf nø t inaf	ke ʃ inaf nø s inaf	ke p inaf nø s inaf	ke ʃ inaf nø t inaf	ke p inaf nø z inaf	ke ʃ inaf nø d inaf
5	novel	ke z afam nø ʒ afam	ke d afam nø b afam	ke d afam nø ʒ afam	ke z afam nø b afam	ke t afam nø ʒ afam	ke s afam nø b afam
6	novel	ke z uve nø ʒ uve	ke d uve nø b uve	ke d uve nø ʒ uve	ke z uve nø b uve	ke t uve nø ʒ uve	ke s uve nø b uve
7	novel	ke z iba nø ʒ iba	ke d iba nø b iba	ke d iba nø ʒ iba	ke z iba nø b iba	ke t iba nø ʒ iba	ke s iba nø b iba
8	novel	ke z inu nø ʒ inu	ke d inu nø b inu	ke d inu nø ʒ inu	ke z inu nø b inu	ke t inu nø ʒ inu	ke s inu nø b inu
9	novel	ke p uda nø t uda	ke ʃ uda nø s uda	ke p uda nø s uda	ke ʃ uda nø t uda	ke p uda nø z uda	ke ʃ uda nø d uda
10	novel	ke p agil nø t agil	ke ʃ agil nø s agil	ke p agil nø s agil	ke ʃ agil nø t agil	ke p agil nø z agil	ke ʃ agil nø d agil
11	novel	ke p iki nø t iki	ke ʃ iki nø s iki	ke p iki nø s iki	ke ʃ iki nø t iki	ke p iki nø z iki	ke ʃ iki nø d iki
12	novel	ke p aʒe nø t aʒe	ke ʃ aʒe nø s aʒe	ke p aʒe nø s aʒe	ke ʃ aʒe nø t aʒe	ke p aʒe nø z aʒe	ke ʃ aʒe nø d aʒe