Perceptual Advantage from Generalized Linguistic Knowledge

Bożena Pająk (bpajak@ling.ucsd.edu)
Department of Linguistics, UC San Diego
9500 Gilman Drive #108, La Jolla, CA 92093 USA

Abstract

We address the question of how previously acquired linguistic knowledge facilitates perception and learning of a new language. We report results from two experiments showing evidence that participants better discriminate a segmental duration contrast in a novel language if they had some previous exposure to a language that uses duration contrastively. Crucially, the perceptual advantage occurs even when the novel language employs the contrast in entirely different conditions: in novel segmental contexts and for novel segments, including a change from application to vowels to application to consonants. We take these results to suggest that language learners use their knowledge of previously learned languages to make inferences about the ways in which languages are likely to vary, which in turn increases their perceptual sensitivity when languages do in fact vary in the predicted ways.

Keywords: Speech perception; language learning; overhypotheses; cross-linguistic influence; multilingualism.

Introduction

Second language (L2) acquisition and bilingualism have received a lot of attention in the literature over the past few decades (for an overview, see e.g., Ritchie & Bhatia, 2009). Significantly less research has been done on acquisition of more than two languages, but it is now more widely recognized that acquisition of additional languages (L*n*) is fundamentally different from L2 acquisition due to more possibilities for between-language interactions (De Angelis, 2007). However, the mechanisms behind these between-language interactions are still very poorly understood.

One specific consequence of this limitation in scope to L2 acquisition has been a lack of systematic research investigating the common intuition that while learning an L2 is often hard, learning each subsequent language becomes easier. We explore this intuition by asking higher-level questions about the learning process, as well as the nature of abstracting and generalizing, with the goal to understand what mechanisms would produce facilitation in L*n* acquisition. A convenient framework for asking these questions makes use of Goodman's (1955) notion of "overhypothesis". Goodman's observation was that humans not only learn things that they experience directly, but they also infer abstract knowledge about how the things they experience directly are structured. Imagine a stack of bags that contain colored marbles. You empty a few of them and discover that some bags only have black marbles, while others only have white ones. Now imagine you pick a new bag and draw one marble which turns out to be white. The experience with previous bags makes you hypothesize that this particular bag contains only white marbles. This hypothesis is based on the overhypothesis you formed through your overall experience with this stack of bags that each bag contains marbles that are uniform in color.

Our proposal is to work within the overhypothesis framework to build a model of multiple language learning (MLL). The reasoning behind the model is that with knowledge of only one language, one possesses a limited amount of information about the possible features that languages can have: for example, what sounds they use or what syntactic and morphological features they employ. For a monolingual speaker, the general overhypothesis over how languages are structured, or what linguistic dimensions are relevant to assign meaning, depends entirely on the knowledge of one's single native language (L1). This means that when learning an L2, one is likely to assume that the L2 features are similar to those of L1, a prediction confirmed by a large body of research in L2 acquisition (Ritchie & Bhatia, 2009). On the other hand, with at least some basic knowledge of two or more languages, one can update the overhypothesis (or, rather, the set of overhypotheses, each related to a specific linguistic dimension) by reevaluating which dimensions are relevant for each language. Structural differences between two or more languages on any given dimension provide a basis for expanding the hypothesis space (i.e., predictions about possible categories) for this dimension. We hypothesize that this conceptually expanded hypothesis space facilitates learning of novel categories along relevant dimensions in an L*n* due to the fact that specific predictions about possible categories have already been formed, which in turn accelerates their processing.

When applied at the level of sound, the model makes predictions regarding the ability of bi/multilingual learners to discriminate novel contrasts along familiar dimensions in an L*n*. Auditory perception of nonnative contrasts is often initially impaired, but can significantly improve with increased exposure. The explanation for this initial difficulty is that nonnative speech perception is shaped by the L1 experience, and current theories are successful in predicting which L2 sounds will be harder to initially perceive by listeners with a given language background (Kuhl & Iverson, 1995; Best, 1995; Flege, 1995). However, these models are not explicit about how the L1 bias is overcome when nonnative phonological categories are successfully learned. A possible answer is instead provided by the literature on perceptual categorization. When learning phonological categories in L2, language learners adjust weights assigned to different phonetic dimensions so that dimensions reliably aiding in proper phoneme categorization in L2 are given more weight, while dimensions creating phonologically irrelevant variation are given almost no weight (Kruschke, 1992; Strange & Shafer, 2008). Incorporating these theoretical assumption, the MLL model predicts that assigning high weight to a given phonetic dimension based on L2 input raises the likelihood of this dimension also being considered as relevant in L*n*. This, in turn, leads to facilitation in perception and learning of novel L*n* categories that make use of this dimension.

As an example, consider the phonetic dimension of segmental duration. Imagine a native speaker of English who also speaks Cantonese. In English, segmental duration is used mainly as a prosodic cue (Klatt, 1976), while in Cantonese, vowel duration can be considered a *contrastive* feature (Bauer & Benedict, 1997), which means that words can potentially be distinguished based solely on the duration of a given vowel (short vs. long). The model assumes that this speaker has formed an overhypothesis over the dimension of duration stating that duration of segments can be relevant for assigning meaning in some languages. Crucially, even though the speaker's experience with duration is only based on *vowel* segments, he/she is expected to have formed an overhypothesis over the duration dimension that is not segment-specific, and consequently, to have formed hypotheses (or predictions) regarding the relevance of duration for *any* segment. Now, if this speaker is learning an L*n* like Polish, which has a *consonant* duration contrast, the model predicts a facilitation in learning this contrast, as compared to a learner who has not had any previous exposure to any duration contrasts, and consequently no opportunity to form an overhypothesis over the duration dimension.

Here we report the results of two experiments, in which the perception of short vs. long consonants was tested. In the first experiment, we tested speakers of American English who had previously learned a language with contrastive consonant duration. We expected to observe a perceptual advantage for this group over English speakers without such experience. Furthermore, we tested whether the perceptual advantage generalizes to novel consonant segments and novel segmental contexts. In the second experiment, we asked the question of whether the feature of contrastive duration can be generalized even further: namely, from vowels to consonants. Specifically, we tested the perception of a consonant duration contrast by speakers fluent in English and Cantonese, which has a duration contrast for vowels, but not for consonants.

Experiment 1

In an AX discrimination task we tested the perception of a consonant duration contrast (short vs. long) by a "bilingual group": native (or near-native) speakers of American English with previous exposure to another language that uses consonant duration contrastively (e.g., Japanese or Italian). The control "monolingual group" consisted of native speakers of American English with no previous exposure to any language that contrasts duration.[1] Following the assumption of the MLL model that the bilinguals have assigned a high weight to the duration dimension in their L2, it was predicted that the bilingual group would perform better than the monolingual group. Furthermore, the perception of the duration contrast was tested in different phonotactic environments (here, the position in a word and the adjacent segments). While some theories assume that learning new contrasts is context-specific (Flege, 1995), the proposed model predicts that the abstracted knowledge should allow generalization across different environments. Additional comparisons within the bilingual group were planned in order to investigate more closely the process of generalization from previous knowledge. In particular, it was predicted that the bilingual participants would be able to generalize their perceptual capacity for duration contrasts to novel segments, and – following the underlying principle of the overhypothesis framework – that familiarity with the contrast in at least two different contexts would faciliate generalization to a novel context more than its familiarity in only one context.

Method

Participants 80 undergraduate students at UC San Diego participated in the experiment for course credit: 40 "monolinguals" and 40 "bilinguals". The bilingual group was largely heterogeneous. It consisted of speakers of a total of 17 different L2s, and varied in their proficiency in L2, as well as the manner of exposure to L2 (school instruction or exposure at home through family members). The bilingual participants were further split in two ways depending on the types of segments possible as long consonants and the positions in which they occur in their L2. The division by "segment" included "[ss] bilinguals" who were only familiar with long [s], and "[ss] & [zz] bilinguals" who were familiar with both long [s] and long [z]. The division by "context" included "intervocalic bilinguals" who were only familiar with long consonants in the intervocalic context, and "intervocalic+ bilinguals" who were familiar with the contrast in the intervocalic context plus in at least one other context (word-medial preconsonantal and/or word-initial prevocalic).

Materials The materials consisted of nonce words constructed around either a long or a short target consonant. The target consonants were placed in four different contexts created by crossing two conditions: word position (medial vs. initial) with following segment (vowel vs. consonant). All the bilingual participants had previous exposure to the contrast in the word-medial prevocalic (or intervocalic) context, while none had previous exposure to the contrast in the word-initial preconsonantal context. Two different types of segments were used: voiceless alveolar fricatives [s]/[ss] and voiced alveolar fricatives [z]/[zz], resulting in a total of eight conditions. The materials are shown in Table 1.

[1] American English does not use duration contrastively. Vowel duration varies, but it correlates with the tense-lax distinction (e.g., *beat* vs. *bit*) or depends on the voicing of the following segment (e.g., *cad* vs. *cat*). Long consonants are sometimes attested but only at morpheme boundaries (e.g., *dissatisfied*; Benus, Smorodinsky, & Gafos, 2003). Minimal pairs are rare (e.g., *unnamed* vs. *unaimed*), and for most speakers the contrast in neutralized (Kaye, 2005). Furthermore, there is evidence that by 18 months of age English-learning infants process duration contrasts differently from infants learning a language that contrasts duration (e.g., Japanese; Mugitani, Pons, Fais, Werker, & Amano, 2008).

Table 1: Materials. (V-vowel, C-consonant, #-word boundary)

		Prevocalic	Preconsonantal
		V_V	**V_C**
Word-medial	*voiceless*	asa/assa	asta/assta
	voiced	aza/azza	azda/azzda
		#_V	**#_C**
Word-initial	*voiceless*	sa/ssa	sta/ssta
	voiced	za/zza	zda/zzda

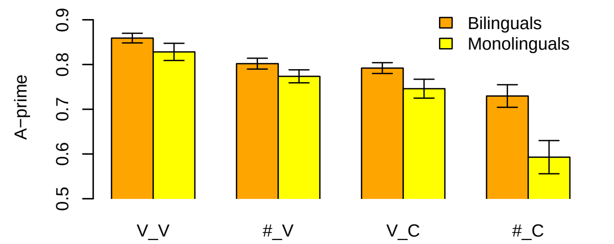

Figure 1: Performance on the short/long consonant contrast in different contexts by monolinguals and bilinguals. (Error bars are standard errors.)

The materials were recorded by a male native speaker of Moroccan Arabic since all the words were phonotactically legal in this language. For each test word, 18 repetitions were recorded. The duration of the fricatives was measured, and five tokens with fricatives that approximated mean duration for each condition were selected for use in the experiment. The trials consisted of an equal number of "different" pairs (e.g., asa-assa) and "same" pairs (e.g., assa-assa or asa-asa). Even in the "same" pairs, the first and second words in each pair were always physically different tokens, and were separated by an interstimulus interval of 500ms. Each participant heard 12 "different" pairs and 12 "same" pairs for each of the eight conditions. There was a total of 384 pairs in the experiment: 192 test pairs and 192 fillers.

Procedure The experiment began with a practice session during which participants listened to 16 filler stimuli (8 "different" and 8 "same" pairs) over headphones, and were asked to respond by clicking on one of two answer boxes displayed on the computer screen saying "same word" or "different words". No feedback was given during the practice session. The test trials followed immediately after the practice session. Participants were presented with six 64-trial blocks. On each trial, a stimulus was presented aurally through headphones, and the participant responded by clicking on one of the two boxes on the computer screen. Each stimulus was played once without a replay option. The response to one stimulus triggered the presentation of the following stimulus with a delay of 500ms. The stimuli order was randomized for every participant. There was a self-terminated break after each block.

Results

We calculated A-prime scores for each participant and each condition (the same results hold for d-prime). A-prime (Grier, 1971) was used to measure the participants' capacity to perceive the short/long consonant contrast, and is a non-parametric analog of d-prime. Both A-prime and d-prime are measures of sensitivity to a given contrast, and are calculated by taking into account the proportion of Hits (responding 'different' when the stimulus is 'different') and False Alarms (responding 'different' when the stimulus is 'same').[2] A-prime yields values between 0 and 1, where 1 means 'perfect discriminability' and 0.5 is chance performance.

Monolinguals vs. bilinguals The results are plotted in Fig. 1. We analyzed the scores using a repeated measures ANOVA with the within-participants factors *position* ('medial' or 'initial'), *following segment* ('vowel' or 'consonant'), *voicing* ('voiced' or 'voiceless'), and the between-participants factor *language background* ('monolingual' or 'bilingual'). There was a significant main effect of *language background* $[F(1,78) = 12.8; p < .001]$ with the bilingual group performing better ($\bar{A'} = 0.79$) than the monolingual group ($\bar{A'} = 0.72$).

There was also a significant interaction between *language background* and *following segment* $[F(1,78) = 11.4; p < .01]$, and a three-way interaction between *language background*, *following segment* and *position* $[F(1,78) = 4.2; p < .05]$. The difference in performance between monolinguals and bilinguals was larger in preconsonantal than in prevocalic contexts, and was especially striking in the word-initial, preconsonantal contexts.

In addition, there were significant main effects of *position* [F(1,78)=34.6; p<.001] and *following segment* $[F(1,78) = 64.8; p < .001]$ independent of language background. That is, both groups performed better when the contrast was in word-medial (vs. word-initial) contexts, and better when it was prevocalic (vs. preconsonantal).

Bilinguals: segments Two groups of bilinguals were compared depending on the segments that occur as long consonants in their L2. The results are plotted in Fig. 2. We analyzed the scores using a repeated measures ANOVA with the same as before within-participants factors *position* ('medial' or 'initial'), *following segment* ('vowel' or 'consonant'), *voicing* ('voiced' or 'voiceless'), and the between-participants factor *L2-long-consonant-segment* ('[ss] bilinguals' or '[ss] & [zz] bilinguals'). The two groups of bilinguals were balanced by randomly removing participants from the larger group, leaving a total of 20 participants for this comparison.

No significant main effect of *L2-long-consonant-segment*

[2] The formula used for calculating A-prime was the following: $A' = 0.5 + \frac{(H-FA)(1+H-FA)}{4H(1-FA)}$ (where H = Hits, and FA = False Alarms; Grier, 1971, p. 425). In order to avoid infinite or undefined values, $H = 0$ was converted to $\frac{1}{2N}$, and $F = 1$ was converted to $1 - \frac{1}{2N}$ (where N = number of trials on which the proportion is based; Macmillan & Creelman, 2005).

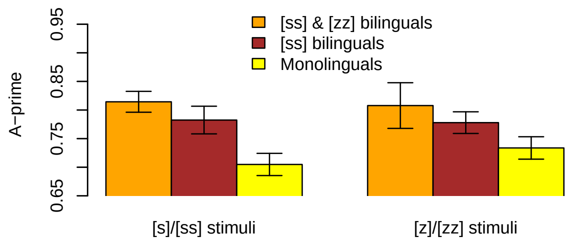

Figure 2: Performance on the short/long consonant contrast by bilinguals familiar with both long [s] and long [z] ('[ss] & [zz] bilinguals') and bilinguals only familiar with long [s] but not long [z] ('[ss] bilinguals'). Monolinguals added for comparison. (Error bars are standard errors.)

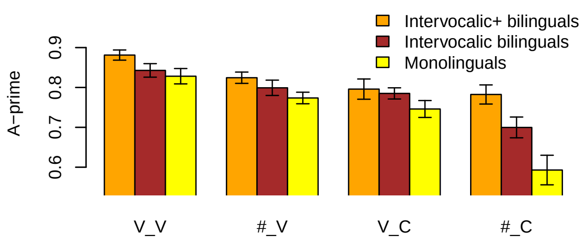

Figure 3: Performance on the short/long consonant contrast by bilinguals familiar with the contrast in different contexts ('intervocalic+ bilinguals') and bilinguals familiar with the contrast only in the intervocalic context ('intervocalic bilinguals'). Monolinguals added for comparison. (Error bars are standard errors.)

was found $[F < 1]$ and no interaction between *L2-long-consonant-segment* and *voicing* $[F < 1]$. Both groups of bilinguals performed the same on both voiceless and voiced contrasts.

There were also significant main effects of *position* $[F(1,18) = 8.1; p < .05]$ and *following segment* $[F(1,18) = 16.3; p < .001]$. Again, both groups performed better when the contrast was in word-medial (vs. word-initial) contexts, and better when it was prevocalic (vs. preconsonantal).

Bilinguals: context The final comparison involved bilinguals which were grouped depending on the contexts in which long consonants are possible in their L2. The results are plotted in Fig. 3. We analyzed the scores using a repeated measures ANOVA with the within-participants factors again being *position* ('medial' or 'initial'), *following segment* ('vowel' or 'consonant'), *voicing* ('voiced' or 'voiceless'), and the between-participants factor *L2-long-consonant-context* ('intervocalic' or 'intervocalic plus V_C and/or #_V'). The two groups of bilinguals were balanced by randomly removing participants from the larger group, leaving a total of 32 participants for this comparison.

There was a significant main effect of *L2-long-consonant-context* $[F(1,30) = 5.5; p < .05]$ with the intervocalic+ bilingual group performing better $(\bar{A}' = 0.79)$ than the intervocalic group $(\bar{A}' = 0.72)$.

As in previous comparisons, there were significant main effects of *position* $[F(1,30) = 18.5; p < .001]$ and *following segment* $[F(1,30) = 19.5; p < .001]$. As before, both groups performed better when the contrast was in word-medial (vs. word-initial) contexts, and better when it was prevocalic (vs. preconsonantal).

Discussion

As predicted, the participants performed better on the short/long consonant contrast than monolingual participants. Importantly, the effect was observed despite high heterogeneity of the bilingual group in terms of their L2 background, the shared feature being the presence of the short/long consonant contrast in every L2. All the bilingual participants seemed to be able to use a similar kind of perceptual capacity which emerged from their different backgrounds. This result provides support for the hypothesis that previous exposure to duration contrasts in an L2 improves perception of a similar contrast in a novel language.

Furthermore, better performance by bilinguals was not simply a result of direct incorporation of certain elements from L2 to a novel language, because – as predicted – the bilinguals were able to generalize their perceptual capacity to novel segments (at least across voicing of segments) and novel contexts. Thus, the perceptual capacity was not found to be context-specific, even though some contexts may be perceptually harder than others.

Finally, following the hypothesis, the bilinguals whose L2 made the contrast in at least two different contexts (intervocalic, word-medial preconsonantal and/or word-initial prevocalic) performed better in the novel word-initial preconsonantal context than the group whose L2 only used the contrast intervocalically. This suggests that, while exposure to at least one segmental context of a duration contrast helps with the overall perception (as is the case for the "intervocalic bilinguals"), it is the exposure to at least two different contexts that allows for a real boost in perceptual capacity (as observed for the "intervocalic+ bilinguals"). This result can be interpreted as a supporting piece of evidence for the overhypothesis framework: exposure to a contrast in at least two contexts allows for the formation of an overhypothesis that this particular contrast can occur in many different contexts. Interestingly, the "intervocalic+ bilinguals" also performed better than the "intervocalic bilinguals" in the intervocalic context, to which all the bilinguals had equal exposure. While this result does not directly follow from the hypothesis, it might be that forming the overhypothesis over contexts makes the perceptual system more attuned to the contrast in any context, either novel or non-novel.

Experiment 2

Experiment 2 was designed in order to test whether the feature of contrastive duration can be generalized further than across voicing of segments, namely, from vowels to consonants. The participants were speakers of Cantonese (also fluent in Mandarin) and speakers of Mandarin with no knowledge of Cantonese. Cantonese has vowel duration contrasts, but no consonant duration contrasts, while Mandarin does not use duration of any segments contrastively. To control for differences in populations, the experiment also included stimuli with a sibilant contrast from Polish, which were chosen because similar consonants form part of the Mandarin consonant inventory. Thus, the two groups of participants (Cantonese/Mandarin and Mandarin) were exposed to two types of stimuli: duration contrasts (short vs. long consonants) and sibilant contrasts (alveolo-palatal vs. postalveolar/retroflex consonants). The MLL model predicted that Cantonese speakers would perform better than Mandarin speakers on the duration contrast due to their experience with the vowel duration contrast in Cantonese. However, both groups were predicted to perform equally well on the sibilant contrast due to their familiarity with a similar contrast in Mandarin (although a slight advantage for the native Mandarin speakers was expected for this contrast).

Method

Participants 40 undergraduate students at UC San Diego participated in the experiment for course credit. 20 were native speakers of Mandarin fluent in English, and the other 20 were native speakers of Cantonese fluent in English and at least competent in Mandarin.

Materials The materials consisted of two types of stimuli, as shown in Table 2.

Table 2: Materials: segmental contrasts.

A. Duration contrasts (short vs. long)

Sonorants	Obstruents
[j]/[jj] [w]/[ww] [l]/[ll] [m]/[mm] [n]/[nn]	[f]/[ff] [s]/[ss]

B. Sibilant contrasts (alveolo-palatal vs. postalveolar/retroflex)

Voiceless	Voiced
[ɕ]/[ʂ] [t͡ɕ]/[t͡ʂ]	[ʑ]/[ʐ] [d͡ʑ]/[d͡ʐ]

This created 4 conditions by crossing two factors: contrast (duration or sibilant) with language background (Cantonese/Mandarin or Mandarin).

Each contrast was embedded in seven different frames: [pa_a], [pe_a], [po_a], [ta_a], [te_a], [ka_a], [ke_a]. All the words were recorded by a phonetically-trained native speaker of Polish with five repetions of each word. One token of each stimulus type was chosen as a frame (a short consonant token for duration contrasts and a postalveolar/retroflex token for sibilant contrasts). The target consonants were spliced out from different tokens. Long consonants were created from the short consonants by either doubling their length (for sonorant consonants: [j], [w], [l], [m], and [n]) or elongating it by half its length (for obstruent consonants: [f] and [s]).[3] The stimuli were created by pairing words that were "different" (e.g., paja-pajja) and "same" (e.g., paja-paja or pajja-pajja). Unlike in Experiment 1, the "same" words in each pair were physically identical and the "different" words in each pair always shared a physically identical frame (i.e., the words were identical except for artificial lengthening for duration contrasts and a spliced consonant for sibilant contrasts). This was done to ensure that any difference in the participants' responses resulted only from the manipulation of interest. The words in each pair were separated by an interstimulus interval of 750ms. Each pair was repeated twice throughout the experiment, which yielded a total of 392 pairs: 196 pairs with duration contrasts and 196 pairs with sibilant or other (filler) contrasts.

[3]This difference was introduced in order to account for the fact that intervocalic duration contrasts are perceptually harder for sonorants than for obstruents (Kawahara, 2007).

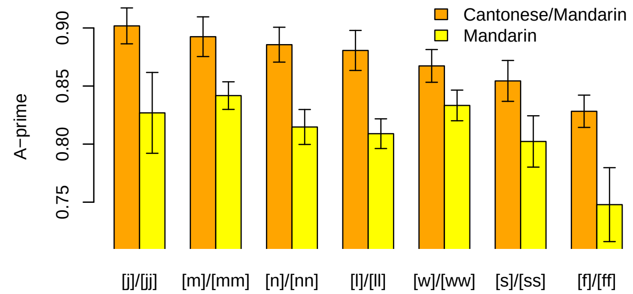

Figure 4: Performance on the long/short consonant contrast by Cantonese/Mandarin and Mandarin speakers. (Error bars are standard errors.)

Procedure The procedure was almost identical to experiment 1. The differences included the number of blocks (seven 56-trial blocks) and the response type: instead of clicking on the screen with a mouse, participants responded by pushing buttons on a game pad.

Results

As in Experiment 1, we calculated A-prime scores for each participant in each condition as a measure of contrast sensitivity (the same results hold for d-prime).

Duration contrast The results from the duration contrasts are plotted in Fig. 4. We analyzed the scores using a repeated measures ANOVA with the within-participants factor *segment* ([j], [w], [l], [m], [n], [f], or [s]), and the between-participants factor *language background* ('Mandarin' or 'Cantonese/Mandarin'). There was a significant main effect of *language background* $[F(1,38) = 12.7; p < .01]$ with the Cantonese/Mandarin group performing better $(\bar{A'} = 0.87)$ than the Mandarin group $(\bar{A'} = 0.81)$. There was also a significant main effects of *segment* $[F(6,228) = 5.8; p < .001]$, indicating that some segments were overall harder than others.

Sibilant contrast The results from the sibilant contrasts are plotted in Fig. 5. We analyzed the scores using a repeated measures ANOVA with the within-participants factor *segment* ([ɕ]/[ʂ], [ʑ]/[ʐ], [t͡ɕ]/[t͡ʂ], [d͡ʑ]/[d͡ʐ].), and the between-participants factor *language background* ('Mandarin' or 'Cantonese/Mandarin'). As predicted for this con-

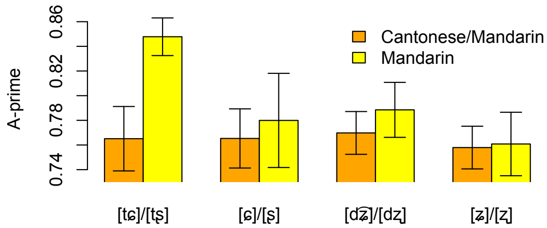

Figure 5: Performance on the sibilant contrast by Cantonese/Mandarin and Mandarin speakers. (Error bars are standard errors.)

trast, no significant main effects of *language background* $[F(1,38) = 1.6; p = .21]$ nor *segment* $[F(3,114) = 1.9; p = .14]$ were found. Both groups performed similarly on all sibilant contrasts, as illustrated in Fig. 5. There was, however, a tendency for Mandarin speakers to perform better than Cantonese speakers, at least on one type of contrast ([t͡ɕ]/[t͡ʂ]).

Discussion

This experiment showed that the speakers of Cantonese/Mandarin perform better on the short/long consonant duration contrast than Mandarin speakers without any exposure to Cantonese. It was hypothesized that such a difference between the two groups would be observed due to the fact that Cantonese uses *vowel* duration contrastively. This result suggests that Cantonese speakers were able to generalize their knowledge about a vowel duration contrast to a consonant duration contrast in a way that perception of the latter contrast was facilitated.

Importantly, the better performance of the Cantonese participants was not due to other differences in populations since the two groups performed equally well on the control contrast of sibilants. In this case, both groups were predicted to perform similarly due to the influence of Mandarin, which has a similar contrast between voiceless sibilants.

The combination of these results means that the feature of contrastive duration can indeed be abstracted away from a limited set of segments (e.g., vowels) and applied in novel conditions with a perceptual advantage, thus supporting the predictions of the model.

Conclusion

This paper argued that previously acquired linguistic knowledge can have a facilitative effect on perception and learning of new languages. In Experiment 1, we showed that participants with previous exposure to a language that uses consonant duration contrastively are better at discriminating a similar duration contrast in a novel language. Perceptual advantage was observed even if the contrast was presented in novel segmental contexts and for novel segments. Experiment 2 showed an even stronger result: perceptual advantage on consonant duration contrasts was observed for participants who only had previous exposure to *vowel* duration contrasts. Together, these results support the MLL model and the overhypothesis framework, which predict that knowledge of previously learned languages is generalized and leads language learners to make inferences about the ways in which languages are likely to vary. These inferences, or overhypotheses, about dimensions along which languages can differ may in turn increase learners' perceptual sensitivity to contrasts that the overhypotheses predict. Having established that such generalization occurs from previously learned to novel languages, the next step for future work will be to determine the exact conditions under which overhypotheses are made, and in what exact ways they are used in language learning.

Acknowledgments

For their helpful feedback on different stages of this project, the author thanks: Abdelhak Akjeje, Eric Baković, Klinton Bicknell, Sarah Creel, Tamar Gollan, Cindy Kilpatrick, Roger Levy, Sharon Rose, and four anonymous reviewers. Christopher Gaudiot and Rachel O'Sullivan helped with data collection for Experiment 1.

References

Bauer, R. S., & Benedict, P. K. (1997). *Modern Cantonese phonology*. Berlin/New York: Mouton de Gruyter.
Benus, S., Smorodinsky, I., & Gafos, A. (2003). Gestural coordination and the distribution of English 'geminates'. In *Proceedings of the 27th Annual Penn Linguistics Colloquium* (pp. 33–46). Philadelphia, PA: University of Pennsylvania.
Best, C. T. (1995). A direct realist view of cross-language speech perception. In W. Strange (Ed.), *Speech perception and linguistic experience: issues in cross-language research* (pp. 171–204). Timonium, MD: York Press.
De Angelis, G. (2007). *Third or additional language acquisition*. Clevedon: Multilingual Matters.
Flege, J. (1995). Second-language speech learning: theory, findings and problems. In W. Strange (Ed.), *Speech perception and linguistic experience: issues in cross-language research* (pp. 229–273). Timonium, MD: York Press.
Goodman, N. (1955). *Fact, fiction, and forecast*. Cambridge, MA: Harvard University Press.
Grier, J. B. (1971). Nonparametric indexes for sensistivity and bias. *Psychological Bulletin*, *75*(6), 424–429.
Kawahara, S. (2007). Sonorancy and geminacy. In *University of Massachusetts Occasional Papers in Linguistics 32: Papers in Optimality III* (pp. 145–186). Amherst: GLSA.
Kaye, A. (2005). Gemination in English. *English Today*, *21*, 43–55.
Klatt, D. H. (1976). Linguistic uses of segmental duration in English: Acoustic and perceptual evidence. *Journal of the Acoustical Society of America*, *59*(5), 1208-1221.
Kruschke, J. K. (1992). ALCOVE: an exemplar-based connectionist model of category learning. *Psychological Review*, *99*(1), 22–44.
Kuhl, P. K., & Iverson, P. (1995). Linguistic experience and the "perceptual magnet effect". In W. Strange (Ed.), *Speech perception and linguistic experience: issues in cross-language research* (pp. 121–154). Timonium, MD: York Press.
Macmillan, N. A., & Creelman, C. D. (2005). *Detection theory: a user's guide (2nd ed.)*. Mahwah, NJ: Lawrence Erlbaum Associates.
Mugitani, R., Pons, F., Fais, L., Werker, J. F., & Amano, S. (2008). Perception of vowel length by Japanese- and English-learning infants. *Developmental Psychology*, *45*(1), 236–247.
Ritchie, W. C., & Bhatia, T. K. (Eds.). (2009). *The new handbook of second language acquisition*. Bingley, UK: Emerald Group Publishing Limited.
Strange, W., & Shafer, V. L. (2008). Speech perception in second language learners: the re-education of selective perception. In J. G. Hansen Edwards & M. L. Zampini (Eds.), *Phonology and second language acquisition* (pp. 153–191). Amsterdam/Philadelphia: John Benjamins.