Journal of Experimental Psychology: Learning, Memory, and Cognition

Why Segmentation Matters: Experience-Driven Segmentation Errors Impair "Morpheme" Learning

Amy S. Finn and Carla L. Hudson Kam

Online First Publication, March 2, 2015. http://dx.doi.org/10.1037/xlm0000114

Journal of Experimental Psychology:
Learning, Memory, and Cognition

0278-7393/15/$12.00 http://dx.doi.org/10.1037/xlm0000114

RESEARCH REPORT

Why Segmentation Matters: Experience-Driven Segmentation Errors Impair "Morpheme" Learning

Amy S. Finn
Massachusetts Institute of Technology

Carla L. Hudson Kam
University of British Columbia

We ask whether an adult learner's knowledge of their native language impedes statistical learning in a new language beyond just word segmentation (as previously shown). In particular, we examine the impact of native-language word-form phonotactics on learners' ability to segment words into their component morphemes and learn phonologically triggered variation of morphemes. We find that learning is impaired when words and component morphemes are structured to conflict with a learner's native-language phonotactic system, but not when native-language phonotactics do not conflict with morpheme boundaries in the artificial language. A learner's native-language knowledge can therefore have a cascading impact affecting word segmentation and the morphological variation that relies upon proper segmentation. These results show that getting word segmentation right early in learning is deeply important for learning other aspects of language, even those (morphology) that are known to pose a great difficulty for adult language learners.

Keywords: language learning, statistical learning, word segmentation, morphology, sensitive period

Languages have statistical structure—regularities in the distribution or placement of sounds, words, kinds or categories of words, and phrases—and recently there has been a great deal of work investigating how learners might use these regularities to acquire language. Most research on this ability, often termed statistical learning, proceeds by asking about the kinds of statistical patterns present in actual languages, then building a small artificial language with the same kind or kinds of patterns, and finally assessing learners' ability to use the pattern to acquire some aspect of the language. For instance, Saffran, Aslin, & Newport (1996) noted, based on Harris (1955), that sequences of sounds that are within a word are more predictable than sequences of sounds that are not within the same word. They tested 8-month-old infants' ability to find words in a continuous speech stream based on these regularities of co-occurrence (in particular, transitional probabilities of syllables), and found that the infants could indeed use this information to segment words out of running speech (Saffran, Aslin et al., 1996). Similar research has since demonstrated learners' abilities with a wide range of linguistically relevant statistical patterns: nonadjacent regularities similar to morphological dependencies (Gómez, 2002) and consonant and vowel tiers (Newport & Aslin, 2004), bimodal distributions similar to phonemic distinctions (Maye, Werker, & Gerken, 2002), and even distributional information relevant for learning grammatical morphology (Finley & Newport, 2010) and forming linguistic categories and syntactic phrases (Reeder, Newport, & Aslin, 2013; Thompson & Newport, 2007; Wilson & Hudson Kam, 2013). There is also evidence that learning does not simply stop when some knowledge state has been reached (cf., Gebhart, Aslin, & Newport, 2009). Recent studies on language processing, for instance, have demonstrated that statistical learning is involved in rapid updating of expectations that influence real-time language processing (Farmer, Monaghan, Misyak, & Christiansen, 2011; Fine, Jaeger, Farmer, & Qian, 2013).

Despite its apparent usefulness for language learning, statistical learning is relatively modality and species independent, that is, many of the things humans can learn in artificial language stimuli they can learn from decidedly nonlinguistic input, and many of the same abilities have been shown in nonhumans species as well (for a review see Romberg & Saffran, 2010). And although statistical learning seems to be relatively age-independent (Saffran, Aslin et al., 1996; Saffran, Johnson, Aslin, & Newport, 1999; Saffran, Newport, & Aslin, 1996), developmental differences in learning are observed.

Generally, these developmental differences appear to be related to emerging knowledge of the learners' first language (L1; Finn & Hudson Kam, 2008; Lew-Williams & Saffran, 2012; Thiessen & Saffran, 2003; Yoshida, Pons, Maye, & Werker, 2010). In the context of infants acquiring their L1, such changes are likely highly beneficial. For older learners trying to acquire a new lan-

Amy S. Finn, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology; Carla L. Hudson Kam, Department of Linguistics, University of British Columbia.

This research was supported by National Institutes of Health Grant HD 048572 to Carla L. Hudson Kam.

Correspondence concerning this article should be addressed to Amy S. Finn, 43 Vassar Street 46-4033, Cambridge, MA 02139. E-mail: amyfinn@mit.edu

guage, however, the changes are likely much less helpful, even to the point of impeding learning. Here we explore that possibility, namely that an adult learner's L1 knowledge can impede statistical learning outcomes for a second language (L2), in the context of morphological learning. In particular, using an artificial language, we examine the impact of L1 phonotactics on adult learners' ability to correctly segment words into their component morphemes; we find that segmentation is impaired when words and component morphemes are structured to conflict with a learner's L1 phonotactic system, but not when L1 phonotactics do not conflict with morpheme boundaries in the artificial language.

The Role of Existing Knowledge in Statistical Learning

Although word segmentation from transitional probability (TP; the probability of an *X-Y* transition given *X*) is robust, studies show that prior experience or knowledge can influence learning in both infants and adults. In infants, emerging experience with language appears to shift how statistical regularities are used. Thiessen and Saffran (2003), for instance, found that while 6-month-olds will use TPs for word segmentation even when they conflict with the characteristic stress pattern of English bisyllabic words (strong–weak), 9-month-olds prefer words consistent with the English pattern such that they will override the TPs. This timing is consistent with earlier work demonstrating that 9-month-old, but not 6-month old, English-learning infants show a preference for words with the strong–weak stress pattern (Jusczyk, Cutler, & Redanz, 1993; Thiessen & Saffran, 2003). It has also been shown that learning sound categories (akin to phonemes) from distributional information takes more exposure after infants have become tuned to the phonemes that are relevant in their native language (Yoshida et al., 2010).

Native language knowledge also influences learning from TPs in adults. In the first experimental demonstration of this, adult native English speakers were exposed to an unsegmented speech stream in which English phonotactics were pitted against TP. When words started with sound combinations that are not permitted word-initially in English, adult learners were unable to use TP information for word segmentation (Finn & Hudson Kam, 2008). L1 influence on segmentation is not limited to English phonotactics. Another study found that adults were similarly unable to use TP information for segmentation when general word-form patterns in Catalan (their native language) were pitted against TP (Toro, Pons, Bion, & Sebastián-Gallés, 2011).

These results make a great deal of sense when viewed from the perspective of the well-known phenomena of interference and transfer in second language acquisition. It has long been known that learners often have difficulty learning aspects of a L2 that differ from those in their L1, and the statistical learning results just discussed begin to provide an explanation for interference (at a mechanistic level): learners have difficulty computing statistical regularities that are in direct conflict with the patterns in their L1. However, late learners' difficulties are not limited to cases where the two languages directly conflict. If it were, English learners should have less difficulty with grammatical gender in Spanish than French-speaking learners, for example, since there is no conflict between grammatical gender assignment in English and Spanish while there sometimes is between French and Spanish. However, research shows that this is not the case (White, Valenzuela, Kozlowska-MacGregor, & Leung, 2004). Recent work using artificial languages is beginning to provide hints about these broader learning difficulties too. Finn, Hudson Kam, Ettlinger, Vytlacil, & D'Esposito (2013) showed that the sound inventory of an L2 affects where in the brain the L2 is processed, something with obvious consequences for learning. And Onnis and Thiessen (2013) showed that L1 regularities can even influence the nature of the computations a learner performs in a novel language. In particular, they found that the dominant word order in a participant's L1 affected whether they computed backward or forward TPs in an artificial language. Thus, despite initial and continuing evidence of widespread statistical learning abilities (in terms of species and domains), more recent work is also beginning to uncover interesting constraints on this kind of learning that arise from what a learner already knows about the domain involved, where domain is construed both narrowly (e.g., word forms) and broadly (e.g., language-wide ordering patterns).

The present study is very much in this same line of work. In particular, we examine whether adults' ability to learn morphological variation, something they are known to have difficulty with (Johnson & Newport, 1989; Newport, 1990), is directly affected by native-language phonotactics.[1] Participants were exposed to a language containing two position-dependent prefixes, each of which had two allomorphs (forms dependent on the phonological context), one for stems beginning in vowels (V) and the other for stems beginning in consonant (C) clusters. The onset clusters (CC) either violated (experimental language) or cohered with (control language) English phonotactics (the native language of the participants). For one of the prefixes, clear segmentation information was provided in the input, but for the other, it was not. Thus, participants had to use distributional information alone to properly segment the second prefix from the roots. If missegmentation due to L1 knowledge bleeds into other aspects of learning, we should see learning in the licit (according to L1 phonotactics), but not the illicit language.

It almost seems self-evident that adult learners will have difficulty with this task. Adult learners are known to have problems with morphological variation, often seeming to prefer a single form when the language has an alternation (e.g., Klein & Perdue, 1993). Indeed, some have suggested that adult learners have difficulty with morphological segmentation more broadly (e.g., Newport, 1990). However, a general problem with alternation in adult learners predicts that *all* learners in our study will have problems learning the allomorphy, and that problems will be evident regardless of whether the morphemes are attached to words that have licit phonotactics (according to their native language) or not. On the other hand, there is reason to predict that no one will have trouble. The input is structured to facilitate learning in two ways. First, the illicit phonotactics are not part of the target morphemes themselves, rather, they are in the word that follows these morphemes. Second, the input is constructed to provide clear segmentation information (via infixes, described below) about all of the word initial clusters in the artificial languages (illicit and otherwise), in

[1] Importantly, the basic ability, using TP information to segment subunits within meaningless words, has already been shown by Finley and Newport (2010) using a much simpler language.

a way that mimics what happens in real languages. Given how the input is structured and the numerous studies showing that adults can learn novel forms in artificial languages via statistical learning, it seems quite possible, indeed even likely, that participants should learn the allomorphic alternations. Our hypothesis, that learners will have difficulty learning allomorphic variation when the words that morphemes are attached to conflict with word-formation rules in their first language (but not otherwise), therefore, is not as obvious an outcome as it might at first seem.

Methods

Participants

Thirty undergraduates (18 females) at the University of California, Berkeley participated for course credit. Participants were native speakers of English who reported normal hearing.

Stimuli and Experimental Manipulation

The overall structure of the artificial languages was designed to mimic particular features of natural languages related to morphological marking, but without the words having any meaning. In particular, the languages had features reminiscent of the grammatical categories nouns and verbs, transitive and intransitive sentences, and case agreement. There were two primary word types (A and B), along with several prefixes and infixes. There were also two sentence types, and the distribution of the words, prefixes, and infixes in the two sentence types created a simulation of morphological and grammatical structure reminiscent of natural languages. Thus, the structure is somewhat more complicated than many artificial languages, but importantly, was the same for both the experimental and control languages.

Table 1 lists the words and affixes used in the experimental and control languages. Given the complex nature of the manipulation (and how it is instantiated in the languages), we explain the word and affix types in some detail before describing the particulars of the sentence structure and stimuli set. All transcriptions of stimuli below (and in figures and tables) are based on English orthography-pronunciation correspondences.

A-words and B-words. The 12 main words in each language were divided into two types, A-words and B-words, with B-words being further divided into subtypes B_1 and B_2. As can be seen in Table 1, there were eight A-words, two B_1-words, and two B_2-words. The two major classes were designed to mimic the grammatical classes noun (A-words) and verb (B-words), and the two kinds of B-words are akin to intransitive and transitive verbs. A-words and B-words differed in three ways. First, their form differed: A-words were bisyllabic, and B-words were monosyllabic. Second, they were different in their distribution: sentences were built around B-words, in that they determined the length of the sentence. In particular, B_1-words occurred in sentences with one A-word (short sentence; like an intransitive sentence), B_2-words occurred in sentences with two A-words (long sentence; like a transitive sentence). Unlike transitive and intransitive sentences in real languages, there were no restrictions on which A- and B-words could go together based on meaning; any A-word could occur with any B-word and vice versa. The order of the words was A–B_1 in short sentences and A–B_2–A in long sentences. The third difference was that the A-words (but not the B-words) were divided into two phonologically defined categories. In each language, four of the A-words begin with a vowel and have either vowel-consonant-consonant-vowel-consonant (VCCVC) or VCCV structure, and the other four begin with a consonant cluster and have either CCVCVC or CCVCV structure. Importantly, these four CC-onset words are illicit (in their onsets) according to English phonotactics in the experimental language, but not the control language. All other words and affixes in the languages are licit according to English phonotactics in both languages.

Prefixes. All A-words were preceded by one of two monosyllabic prefixes. These prefixes are the target morphemes, the morphemes that evince allomorphy.

Both prefixes had two forms, or allomorphs: *juh* or *juhb* and *woy* or *woys*. Which form the prefix took depended crucially on the A-word it preceded. Specifically, if the A-word started with a vowel (e.g., *ahbkag*), the prefix would end in a consonant (*juh***b** or *woy***s**). If the A-word started with a consonant cluster (as in *krahrey*), the prefix would take the shorter form and end in a vowel (*juh* or *woy*). This form-contingent structure was modeled on allomorphs that occur in natural languages.[2] This did not differ between the two languages. The relationship between the first sound in the A-word and the form of the prefix was the same in the licit and illicit cluster languages. Indeed, the prefixes themselves were the same in the two languages.

The two prefixes had different distributions. The first prefix (P_1; *juh*) was always attached to the first A-word in the sentence. This prefix is like subject marking or nominative case. The second prefix (P_2; *woy*) was attached to the second A-word in the sentence, and therefore occurred only in long sentences. This prefix is like object or accusative case marking.[3]

Infixes. Finally, there were four monosyllabic infixes that occurred between P_1 and the A-word in short (intransitive) sentences. Like the prefixes, the infixes were contextually dependent in that they did not occur with P_2 or in long sentences. For instance, the infix *muwj* would occur between *juhb* and *ahbkag* (as in *juhbmuwjahbkag*) or between *juh* and *krahrey* (as in *juhmuwjkrahrey*), likewise for every other P_1 and A-word combination. These infixes clearly indicate the boundary between the prefix and the following A-word, and in so doing, stack the deck against our hypothesis. In terms of transitional probability, the infix creates an additional nonadjacent association between the A-word and the prefix (it is additional because the nonadjacent relationship only applies when the infix is present; when it is absent, the relationship is an adjacent one). Importantly, nonadjacent relationships based on TPs are learnable (Gómez, 2002; Newport & Aslin, 2004). The infixes also make the language more naturalistic; in natural languages, forms like nouns are surrounded by multiple different morphemes, each of which potentially serves to assist a learner in

[2] Although there is evidence that suffixes might be easier to learn than prefixes (e.g., (Slobin, 1973; St Clair, Monaghan, & Ramscar, 2009), we used prefixes because it allowed us to use clusters that have previously been shown to cause adult learners difficulty in a word segmentation study (Finn & Hudson Kam, 2008).

[3] Although the forms modelled on morphemes have no associated meanings, distributionally, they are very much like grammatical morphemes, in that their occurrence is correlated with the placement or location of the word they are part of or attached to.

Table 1
Words, Prefixes, and Infixes in the Artificial Languages (Experimental and Control)

Word type				
A (Experimental)	A (Control)	B (Both)	Prefix (Both)	Infix (Both)
chpughzeen	thrughzeen	hahn	juh/juhb	muwj
tfohbuwl	zwohbuwl	chiyd	woy/woys	nuwl
thmahrey	krahrey	geel		reeg
vteesah	bleesah	nahg		neyd
ahbkahg	ahbkahg			
eefgeel	eefgeel			
aevnah	aevnah			
ohnlew	ohnlew			

finding the boundaries (e.g., plural, possessive, agentive, etc.) All four infixes occurred an equal number of times in the input stimuli and equally often between all P_1 and A-word pairings.

Sentences. As shown in Figure 1, short sentences always begin with a P_1 (*juh* or *juhb*), which is followed by an infix, then an A-word, and then a B_1-word. Example short sentences are shown in Figure 1 for both the experimental and control languages. As was described previously, the phonological form of the A-word determines which form of the prefix is used, either *juh* or *juhb*.[4] Long sentences also begin with P_1, either *juh* or *juhb,* followed by an A-word, then a B_2-word, then P_2 (*woy* or *woys*), and then finally the second A-word. The phonological form of the A-word again determines which form of the prefix is used, either *juh* or *juhb* for the first A-word, and either *woy* or *woys* for the second A-word.

Sentences in the presentation set were generated randomly, and then the generated sentences were presented randomly (the same random list for all participants) with the constraint that a sentence could not follow itself. All words were matched for frequency across sentence types (i.e., all eight A-words occurred equally often in short and long sentences). In long sentences, any given A-word occurred equally often as the first and second A-word. In total, the exposure contained 4,480 sentences. Each A-word occurred 840 times, each B-word 1,120 times, each infix 560 times. P_1 occurred 4,480 times, 2,240 times in short sentences and 2,240 times in long sentences, and P_2 occurred 2,240 times (in long sentences only). The exposure set lasted 114 min, which was split up into two sets of stimuli. The duration of the first set was 60 min, and the second set, 54 min.

Short sentence

Structure

$P1_{\text{Dependent on subsequent A}}$ – Infix – $A_{\text{Class 1or2}}$ – B_1

Experimental

*juh—muwj—*thmahrey*—hahn*

Control

juh—muwj—krahrey—hahn

Long sentence

Structure

$P1_{\text{Dependent on subsequent A}}$ – $A_{\text{Class 1or2}}$ – B_2 – $P2_{\text{Dependent on subsequent A}}$ – $A_{\text{Class 1or2}}$

Experimental

*juhb—ahbkag—hahn—woy—*thmahrey

Control

juhb—ahbkag—hahn—woy—krahrey

Figure 1. Structure of the artificial language.

The result of this structure is perfect syllable-to-syllable TPs within words, but not across word or morpheme boundaries. At the level of phonemes, the TPs between the A-words and their prefixes are likewise higher within words than across boundaries. Notably, the fact that A-words can occur after both prefixes and the infix should make them segmentable as units even for the CC-onset words in the experimental language. Indeed, the infix that occurs between P_1 and the A-word in short sentences should make the boundary especially clear.

The auditory stimuli were created with a text-to-speech program SoftVoice (Katz, 2005), which was programmed to produce syllables with a F_0 (fundamental frequency) of 83.62 Hz, to match vowels for length, and to have no coarticulation to avoid non-TP cues to segmentation. Sentences were separated by a pause of 87–93 ms, but there were no pauses between the words within a sentence.

Exposure

Participants were randomly divided into two exposure conditions: experimental and control. In both, participants were told that they were going to listen to a new artificial language, and they were instructed to neither overthink nor ignore the stimulus. To facilitate this, participants colored (using crayons or markers) while they were listening. Exposure lasted for 60 min on the first day and 54 on the second and always occurred on successive days. After exposure on the second day, participants completed a forced-choice test.

Tests

The tests examined whether participants learned that each of the prefixes should end in a consonant if it occurs before an A-word that starts with a vowel and, conversely, should end in a vowel if it occurs before an A-word that starts with a consonant cluster (i.e., juhahbkahg is not grammatical while juh***b***ahbkahg is and vice versa for words that begin with a consonant).

[4] Note that the form of the prefix does not vary with the infix, as infixes always have CVC form. Thus, one can think of the word forms as being determined in stages, with the infix “added after” the prefix form is determined.

Familiar-infix test items. Since we were specifically interested in exactly how participants had segmented the stimuli at the level of the phoneme, we needed a test that enabled us to examine very small differences in placement of the boundary, which the use of a test infix allowed. (Pilot work revealed that participants were distracted by a novel infix at testing if one had not previously occurred during exposure, hence the use of the infixes in the exposure stimuli.) Thus, we tested segmentation by putting a familiar infix, that is, one experienced in the input, within a prefix + A-word string in either the correct or incorrect location with respect to the correct segmentation. Recall that during exposure, infixes had occurred with the first A-word in the sentence, but not the second. Thus, participants had only experienced P_1, but not P_2, in a way that clearly indicated the correct segmentation of A-words.

For half of the forced-choice test items, the initial syllable of the A-word started with a consonant cluster (CCV; eight test items), and the other half started with a vowel (VC; eight test items). Within each initial syllable type (CCV and VC), half of the items used the prefix (P_1; juh) that occurred with the infixes in the training set, and the other half tested items using P_2 (woy). There were a total of eight test items for each prefix, four with CCV-initial syllables, and four with VC-initial syllables. For test items using A-words that began with consonant clusters (CCV-initial syllable items), the infix was in either the correct or incorrect location, and when in the incorrect location, the cluster was split (e.g., juh*infix**ch***pughzeen vs. juh***ch****infix*pughzeen; *juhch* would be incorrect based on the input statistics but would sound better according to English phonotactics).

For VC-initial-syllable test items, there were two types of test items (four each, two using P_1 and two using P_2). In one, the infix was in the correct versus incorrect location with respect to the consonant that is part of the prefix (i.e., juh***b****infix*ahbkahg vs. juh*infix***b**ahkahg; juh***b****infix*ahbkahg would be correct based on the input statistics). These are called VC-prefix-consonant-incorrect-location test items. In the other, the final consonant of the prefix was either present and in the correct location (as it should be) or missing (i.e., juh***b****infix*ahbkahg vs. juh*infix*ahkahg; juh***b****infix*ahbkahg would be correct based on the input statistics). These are called VC-prefix-consonant-missing items. The first type assessed whether participants knew which sounds the prefix consonant "goes with." The latter assessed whether they knew the distributional restrictions on the CV form of the prefix. Note that neither the correct nor incorrect options violate English phonotactic restrictions. These variations in incorrect forms for the VC-initial-syllable A-words were used because it was not clear exactly how the phonotactic restrictions inherent in the CCV-initial-syllable A-words would affect participants' ability to learn the much easier allomorph for the VC-initial-syllable words (if indeed they did), and we did not want to select the wrong possibility and miss finding an effect. Examples of each of these test items can be seen in Figure 2.

Novel-infix test items. We were interested in assessing the strength of any morpheme knowledge acquired during exposure, and so we included eight additional test items that used two novel infixes (*juwg* and *leerg*), four items with P_1, and four with P_2. Half had A-words with CCV-initial syllables, and half had VC-initial syllables. Because of the small amount of items, all VC initial syllable items were of the VC-prefix-consonant-incorrect-location type. These test items examined how robust participants' knowledge was by probing how they would deal with novel forms interacting with the artificial language morphemes.

CCV-initial syllable items

juh*infix***ch**pughzeen vs. juh**ch***infix*pughzeen

VC-initial-syllable test items

juh**b***infix*ahbkahg vs. juh*infix***b**ahkahg
or
juh**b***infix*ahbkahg vs. juh*infix*ahkahg

Figure 2. Example test items.

There were a total of 24 test items (16 using familiar infixes and eight using novel infixes), presented in random order. For each subtest, the correct answer appeared in the first position (of the two presented alternatives) 50% of the time and the second position 50% of the time. Individual test items were the same for each participant. Test items were presented over headphones, with a 1-s pause in between items, and participants were asked to indicate with a button press which member of the pair sounded more like it belonged in the language they were listening to, by pressing "1" if they thought the first option was better and "2" if they thought the second option was better.

Procedure

On the first day, participants came into the lab, read and signed the consent form, and then listened to the language for 60 min. On the second day, they listened to the remaining exposure (54 min) and then performed the tests. Exposure and testing were conducted individually in a quiet room. Stimuli and tests were presented via noise-cancelling headphones. After completing the tests, participants filled in a survey probing their demographic and language backgrounds.

Predictions

If participants are segmenting solely on the basis of distributional information, and adults have no difficulty with such analyses, we would expect them to be able to learn the two different forms for each of the prefixes, as well as how they are distributed, that is, that the CVC form goes with vowel-initial A-words and the CV form with CC-initial words. If native language phonotactic restrictions affect learning, however, we would expect successful learning of the allomorphic alternations only for the control group. The experimental group by contrast will have difficulty. We expect the experimental group to only have difficulty with P_2 given that P_1 segmentation was made clear by the inclusion of the infix in the short sentences, which broke the P_1-to-A-word strings into their component parts and created a boundary. We also predict that difficulty with P_2 would be observed only for the CV allomorph and the CC initial words; the boundary for the CVC allomorph is unaffected by the L1 phonotactic restrictions. Thus, the CVC allomorphs are a sort of language internal control. Additionally, they serve as possible models for (mis)parsing the CV + CC initial words; they could lead participants to extract a CVC form as the basic prefix type, which would lead to misparsing in the case of the CV allomorph.

A third possibility is suggested by the not uncommon belief that adult learners initially extract large chunks from input, due to their superior memory skills, only progressing to analyze the internal structure of the chunks with continued exposure (see, for instance, a thread on idiom comprehension on the CHILDES listserve in early mid-June, 2014). Although there is only mixed support for this idea (see Ellis, 2012 and Weinert, 1995, for reviews), it would predict that neither group would do very well at learning the allomorphy. For a learner to track distributions over words or parts of words (Thompson & Newport, 2007), the units need to be segmented (and represented) correctly. Logic dictates that if they are missegmented early on, learning distributional information about how the units operate will be problematic, especially for subtle aspects of morphology such as allophony, where the form of a morpheme depends on the form of the stem to which it is attached.

Results

Familiar-Infix Test Items

Mean performance on the familiar-infix test items is shown in Figure 3. Our primary hypothesis was that the experimental group would perform poorly (in general and relative to the control group) on CCV-initial-syllable items (that violate English-onset phonotactics) for the prefix that was not segmented in the exposure: P_2. In line with this, a repeated measures analysis of variance (ANOVA) revealed a significant prefix by condition interaction, $F(1, 28) = 4.225$, $p = .049$, $n_p^2 = .131$, stemming from the fact that the simple main effect of prefix was significant for the experimental, $t(14) = 3.228$, $p = .006$, $d = .836$, Bonferroni corrected $\alpha = .025$, but not control, $t(14) = .764$, $p = .458$, $d = .204$, participants. That is, as can be seen in Figure 3a, only the experimental participants had more difficulty with the prefix for which segmentation had to be accomplished via TP information. Moreover, while the control participants' performance was better than chance for both prefixes, P_1: $t(14) = 2.47$, $p = .027$, $d = .637$; P_2: $t(14) = 2.43$, $p = .029$, $d = .628$, participants in the experimental condition only performed better than chance for P_1, $t(14) = 2.75$, $p = .016$, $d = .710$, but not P_2, $t(14) = -1.435$, $p = .173$, $d = .370$. Since the experimental group differs from the control only in whether the CC onset clusters violate (or not) English onset phonotactics, the data indicate that the learning of morphological variation can be specifically affected by native language phonological knowledge.

The VC-initial-syllable test items (shown in Figure 3b) served as a built-in control. Because these words begin with vowels, no word-onset phonotactics are being violated in English. Thus performance should not differ across the groups, and prefix segmentation should not suffer. Recall that we included two different types of incorrect forms in this test, leading to an analysis with two within-subjects (prefix, type) and 1 between-subjects factor (group/condition). The ANOVA revealed no group by prefix, $F(1, 28) = .130$, $p = .721$, $n_p^2 = .005$, or group by type, $F(1, 28) = .207$, $p = .653$, $n_p^2 = .007$, interactions. As predicted, the groups did not treat these items differently. Of interest, however, is a significant prefix by test type interaction, $F(1, 28) = 16.129$, $p < .001$, $n_p^2 = .365$, such that for P_1 and P_2, performance did not differ when the consonant was absent, $t(29) = 0.00$, $p = 1.0$, $d = 0$, but did when it was misplaced, $t(29) = 4.822$, $p < .001$, $d = .886$, Bonferroni corrected $\alpha = .025$. For the unsegmented prefix therefore, all learners performed worse when the letter was present than when it was absent. In fact, performance was significantly below chance specifically when P_2 was present and in the wrong location, $t(29) = -5.113$, $p < .001$, $d = .934$, but significantly above chance for all other subtests (All VC; P_1-present: $t(29) = 3.067$, $p = .005$, $d = .560$; P_1-absent: $t(29) = 2.504$, $p = .018$, $d = .457$; P_2-absent: $t(29) = 2.283$, $p = .030$, $d = .417$). This suggests that learners consistently want to put the letter in the wrong place (in front of the A-word instead of at the end of the prefix). Because learners had not previously been given a segmentation cue for these items, in contrast to P_1, it seems a preference for words to start with consonants is overriding the statistical regularities that point to the consonant being at the end of the prefix.

Novel-Infix Test Items

As with the familiar test items above, novel test items were analyzed separately for CCV-initial-syllable and VC-initial-syllable items, since segmentation should be interrupted only for CCV-initial items (that violate English onset phonotactics in the experimental condition) for the second prefix (P_2) in experimental

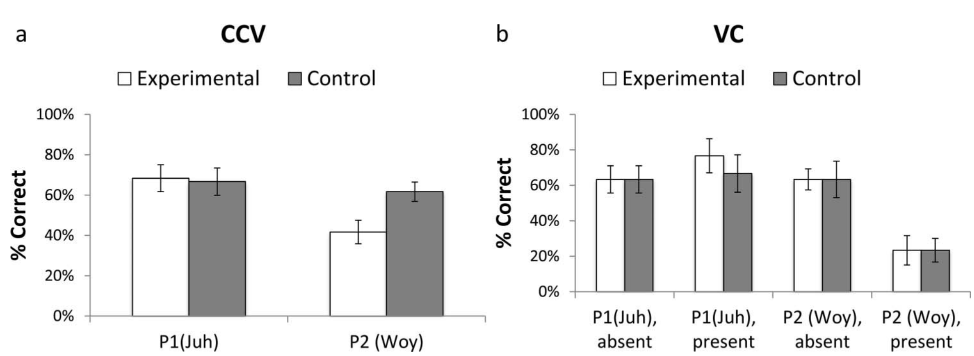

Figure 3. Performance on segmentation tasks for CCV-syllable-initial (a) and VC-syllable-initial (b) A-words, by condition and prefix (and test type). In all cases, error bars represent standard error of the mean.

subjects. However, a repeated measures ANOVA on CCV-initial-syllable items revealed no main effect of group, $F(1, 28) = .974$, $p = .332$, $\eta_p^2 = .034$, or prefix type (P_1 or P_2), $F(1, 28) = .586$, $p = .450$, $\eta_p^2 = .021$, and no prefix by group interaction, $F(1, 28) = 0$, $p = 1.0$, $\eta_p^2 = 0$ (Figure 4a). Thus, experimental and control participants did not differ in their segmentation of these items. One-sample *t* tests revealed that neither experimental, nor control subjects were significantly different from chance on the CCV-initial-syllable-novel items (experimental: P_1: $t(14) = 0$, $p = 1.0$, $d = 0$; P_2: $t(14) = .807$, $p = .433$, $d = .208$; control: P_1: $t(14) = .807$, $p = .433$, $d = .208$; P_2: $t(14) = 1.74$, $p = .104$, $d = .449$).

Recall that we did not predict group differences for VC-initial-syllable items. Unlike familiar infix test items (presented above), there is only one subtype of VC-initial-syllable test items: for both P_1 and P_2 items, the letter in the foil was always present. A repeated measures ANOVA on VC-initial-syllable items revealed no significant effects: group, $F(1, 28) = 4.142$, $p = .051$, $\eta_p^2 = .129$, prefix type, $F(1, 28) = .000$, $p = 1.0$, $\eta_p^2 = 0$, prefix by group interaction, $F(1, 28) = 0$, $p = 1.0$, $\eta_p^2 = 0$ (Figure 4b). Control subjects were marginally better than experimental subjects on this test type. As with the CCV items, one-sample *t* tests revealed that neither experimental, nor control subjects were significantly different from chance on the VC-initial-syllable-novel items (experimental: P_1: $t(14) = -.807$, $p = .433$, $d = .208$; P_2: $t(14) = -1.468$, $p = .164$, $d = .379$; control: P_1: $t(14) = 1.146$, $p = .271$, $d = .296$; P_2: $t(14) = 1.382$, $p = .189$, $d = .357$). Overall, neither the control nor the experimental group segmented in a consistent way for either prefix when novel infixes were used during test. This result was anticipated given pilot work, but we were interested to see if providing a context for the first prefix might make novelty easier for at least the P_1 across the groups of learners. This was not the case.

Discussion

In this article, we ask whether missegmentation in adult learners resulting from a conflict between their L1 and a new language could lead to mislearning of important morphological variation. We exposed learners to one of two languages: control and experimental. While both languages contained consonant clusters, only those in the experimental language violated English phonotactics. As predicted, we found that participants learned the morphological variation in the control language. This was not true for learners in the experimental language—they had difficulty learning the allomorphy, although the statistics in the input clearly indicated where the morpheme boundaries should be. As such, prior linguistic knowledge can have an impact on distributional learning beyond word segmentation.

Two controls were built into this study. First, there is a full control language (learned by a separate group). This language contains all of the complexity of the experimental language (in terms of patterns of words and morpheme structure) but complies with the rules of English word-onset phonotactics. As such, simple complexity cannot be the reason that experimental subjects were unable to learn P_2 (the woy allomorph); rather, the failure of the experimental group to learn this has to do with missegmentation due to illicit onset phonotactics.

A second control was built in by having two, rather than just one, morphemes that vary depending on the word they proceed: P_1 and P_2. Exposure to P_1 was very different from exposure to P_2. During exposure to P_2, learners never heard anything in between the prefix and the following word. Therefore, learning that woy precedes cluster-initial words and woys precedes vowel initial words could be disrupted by missegmenting the cluster and inappropriately appending the initial consonant of the cluster to the end of the prefix. In the case of P_1, the presence of infixes essentially performed this segmentation for learners. Learners hear *juhb-infix-vowel initial word* and *juh-infix-consonant cluster-initial word* and can therefore decipher the onset cluster of any given word from the end of the infix. It is compelling that performance is selectively impaired for learning only the one morpheme, and that, only in the experimental, but not control, language.

These data lead to important questions about how and why knowledge of one's native language influences segmentation in the first (Finn & Hudson Kam, 2008). This experiment, along with previous work, can only document the influence of this knowledge, but does not speak to how exactly the L1 knowledge leads to the effects seen here. It is as yet unknown whether such knowledge interferes with the ability to compute TPs or whether the TPs are computed but when endorsing word candidates it is ignored or overruled by other (and older) knowledge. If L1 knowledge interferes with the ability to compute the TPs, how might this be so? Might learners be drawing associations across the wrong units (or representations), incorrectly calculating associations, or an interaction between the two that leads to the learning patterns (and

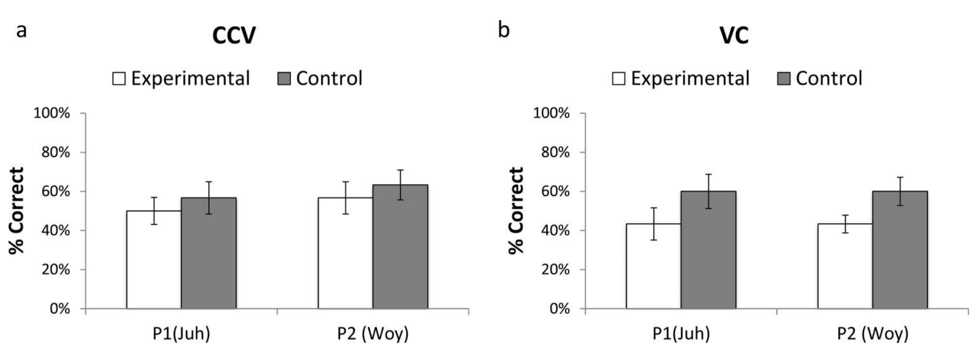

Figure 4. Performance on segmentation task for novel infix for CCV-syllable-initial (a) and VC-syllable-initial (b) A-words, by condition and prefix.

failures) we see? Is it the case, for example, that L2 learners cannot create the appropriate representations to track distributional information over? Or might it be possible (though difficult) to create the appropriate representations, but performing computations over those representations is impaired (perhaps because forming the representations is difficult)? Sorting out these possibilities must be the topic of future work. Demonstrations of L1 transfer or interference in experiments like this establishes artificial language studies as a viable paradigm for use in probing L1 effects in L2 learning more directly. Indeed, the phenomenon of interference has been known to exist for a very long time (Odlin, 1989), but exactly what causes interference effects is difficult to uncover in studies of naturalistic acquisition. Moreover, using very controlled L2s as we have here has the potential to uncover L2 effects that are much broader than those that people usually think about when they think of interference (as in, beyond the knowledge of one linguistic feature in an L1 interfering with learning that feature in an L2, see also Finn et al., 2013).

Two other aspects of our data deserve comment. The first is the asymmetrical results in the VC-initial-syllable items. Recall that all participants performed poorly on the test items where the incorrect options had the final prefix consonant attached to the beginning of the A-word. This finding is based on a small number of test items, and so we wish to be very cautious about interpreting this finding. However, it does suggest that researchers need to be mindful of the predispositions that participants bring to experiments when designing (and so interpreting) test items. (It is not lost on us that this is actually our main point—that participants are not blank slates.) Had we only included this kind of test item, we would have erroneously concluded that participants could not segment those kinds of items at all. Where this preference comes from is worth investigating further, sorting out whether this is a universal preference regarding word forms or is instead experience-based, for instance. Second is the poor performance on the items using novel infixes. These were included to assess whether the knowledge participants ended up with was robust enough to support extension to novel forms. Although participants were clearly able to segment the morphemes and learn the generalizations about their distribution (in the control condition and for one of the prefixes in the experimental condition), they were unable to use what they knew to perform well on the items containing novel infixes. There are two possible explanations for their failure on these items. One is that their knowledge was simply not robust enough to allow them to deal with novel forms in addition to the artificial language morphemes they had been exposed to: exposure was longer than in many artificial language studies, but it was still only two sessions. If this is the case, changing the experiment so that their knowledge is more robust, for example, by increasing exposure time, or making the forms meaningful (see, e.g., Arnon & Ramscar, 2012) might increase performance on these novel test items. Alternatively, they may have had difficulty with the novel forms themselves. Hearing new things in the context of the study might have thrown participants off. If so, then familiarizing participants with novel forms (but not the grammatical combinations to-be-tested) before testing would improve performance.

With these data, we join a growing literature showing that learning in one linguistic domain influences learning in another. Instead of previous reports showing that children's learning in one domain facilitates another (Kuhl, Conboy, Padden, Nelson, & Pruitt, 2005; Newman, Ratner, Jusczyk, Jusczyk, & Dow, 2006; Thal, Bates, Zappia, & Oroz, 1996), we show that mislearning in one domain disrupts learning in another. Thus, our data are in line with other reports showing that one must have the base (or subordinate) units rights to compute the relations among those units (Emberson, Liu, & Zevin, 2013). However, we do not want to suggest that these relationships are deterministic: it is entirely possible that an early missegmentation could later be overcome with additional exposure, a different suite of segmentation cues, and explicit teaching (something typical in adult language acquisition). Ultimately, these possibilities require further research.

The data reported here are clear: adults' native language knowledge can interfere with the learning of important morphological variation in a novel unsegmented language. These data are not documenting the influence of native language in the classic sense: with knowledge of one feature of language—say morphology—interfering with the learning of that same feature in a new language (morphology). Instead, we show that knowledge of native language phonotactics can have a cascading impact affecting word segmentation and the morphological variation that relies upon proper segmentation.

References

Arnon, I., & Ramscar, M. (2012). Granularity and the acquisition of grammatical gender: How order-of-acquisition affects what gets learned. *Cognition, 122,* 292–305. http://dx.doi.org/10.1016/j.cognition.2011.10.009

Ellis, N. C. (2012). Formulaic language and second language acquisition: Zipf and the phrasal teddy bear. *Annual Review of Applied Linguistics, 32,* 17–44. http://dx.doi.org/10.1017/S0267190512000025

Emberson, L. L., Liu, R., & Zevin, J. D. (2013). Is statistical learning constrained by lower level perceptual organization? *Cognition, 128,* 82–102. http://dx.doi.org/10.1016/j.cognition.2012.12.006

Farmer, T. A., Monaghan, P., Misyak, J. B., & Christiansen, M. H. (2011). Phonological typicality influences sentence processing in predictive contexts: Reply to Staub, Grant, Clifton, and Rayner (2009). *Journal of Experimental Psychology: Learning, Memory, and Cognition, 37,* 1318–1325. http://dx.doi.org/10.1037/a0023063

Fine, A. B., Jaeger, T. F., Farmer, T. A., & Qian, T. (2013). Rapid expectation adaptation during syntactic comprehension. *PLoS ONE, 8,* e77661. http://dx.doi.org/10.1371/journal.pone.0077661

Finley, S., & Newport, E. L. (2010). *Morpheme segmentation from distributional information.* Paper presented at the Boston University Conference on Language Development (BUCLD), Boston. MA.

Finn, A. S., & Hudson Kam, C. L. (2008). The curse of knowledge: First language knowledge impairs adult learners' use of novel statistics for word segmentation. *Cognition, 108,* 477–499. http://dx.doi.org/10.1016/j.cognition.2008.04.002

Finn, A. S., Hudson Kam, C. L., Ettlinger, M., Vytlacil, J., & D'Esposito, M. (2013). Learning language with the wrong neural scaffolding: The cost of neural commitment to sounds. *Frontiers in Systems Neuroscience, 7,* article 85. http://dx.doi.org/10.3389/fnsys.2013.00085

Gebhart, A. L., Aslin, R. N., & Newport, E. L. (2009). Changing structures in midstream: Learning along the statistical garden path. *Cognitive Science, 33,* 1087–1116. http://dx.doi.org/10.1111/j.1551-6709.2009.01041.x

Gómez, R. L. (2002). Variability and detection of invariant structure. *Psychological Science, 13,* 431–436. http://dx.doi.org/10.1111/1467-9280.00476

Harris, Z. (1955). From phoneme to morpheme. *Language, 31,* 190–222.

Johnson, J. S., & Newport, E. L. (1989). Critical period effects in second language learning: The influence of maturational state on the acquisition of English as a second language. *Cognitive Psychology, 21,* 60–99. http://dx.doi.org/10.1016/0010-0285(89)90003-0

Jusczyk, P. W., Cutler, A., & Redanz, N. J. (1993). Infants' preference for the predominant stress patterns of English words. *Child Development, 64,* 675–687. http://dx.doi.org/10.2307/1131210

Katz, J. (2005). SoftVoice (Demo Program) [Computer Software]. Los Angeles, CA: SoftVoice, Inc.

Klein, W., & Perdue, C. (1993). Utterance structure. In C. Perdue (Ed.), *Adult language acquisition: Cross-linguistic perspectives. Vol. 2. The results* (pp. 3–40). New York, NY: Cambridge University Press.

Kuhl, P. K., Conboy, B. T., Padden, D., Nelson, T., & Pruitt, J. (2005). Early speech perception and later language development: Implications for the "critical period." *Language Learning and Development, 1,* 237–264. http://dx.doi.org/10.1080/15475441.2005.9671948

Lew-Williams, C., & Saffran, J. R. (2012). All words are not created equal: Expectations about word length guide infant statistical learning. *Cognition, 122,* 241–246. http://dx.doi.org/10.1016/j.cognition.2011.10.007

Maye, J., Werker, J. F., & Gerken, L. (2002). Infant sensitivity to distributional information can affect phonetic discrimination. *Cognition, 82,* B101–B111. http://dx.doi.org/10.1016/S0010-0277(01)00157-3

Newman, R., Ratner, N. B., Jusczyk, A. M., Jusczyk, P. W., & Dow, K. A. (2006). Infants' early ability to segment the conversational speech signal predicts later language development: A retrospective analysis. *Developmental Psychology, 42,* 643–655. http://dx.doi.org/10.1037/0012-1649.42.4.643

Newport, E. (1990). Maturational constraints on language learning. *Cognitive Science, 14,* 11–28. http://dx.doi.org/10.1207/s15516709cog1401_2

Newport, E., & Aslin, R. N. (2004). Learning at a distance: I. Statistical learning of non-adjacent dependencies. *Cognitive Psychology, 48,* 127–162. http://dx.doi.org/10.1016/S0010-0285(03)00128-2

Odlin, T. (1989). *Language transfer: Cross-linguistic influence in language learning*. New York, NY: Cambridge University Press. http://dx.doi.org/10.1017/CBO9781139524537

Onnis, L., & Thiessen, E. (2013). Language experience changes subsequent learning. *Cognition, 126,* 168–284.

Reeder, P. A., Newport, E. L., & Aslin, R. N. (2013). From shared contexts to syntactic categories: The role of distributional information in learning linguistic form-classes. *Cognitive Psychology, 66,* 30–54. http://dx.doi.org/10.1016/j.cogpsych.2012.09.001

Romberg, A. R., & Saffran, J. R. (2010). Statistical learning and language acquisition. *Wiley Interdisciplinary Reviews, 1,* 906–914. http://dx.doi.org/10.1002/wcs.78

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science, 274,* 1926–1928. http://dx.doi.org/10.1126/science.274.5294.1926

Saffran, J. R., Johnson, E. K., Aslin, R. N., & Newport, E. L. (1999). Statistical learning of tone sequences by human infants and adults. *Cognition, 70,* 27–52. http://dx.doi.org/10.1016/S0010-0277(98)00075-4

Saffran, J. R., Newport, E. L., & Aslin, R. N. (1996). Word segmentation: The role of distributional cues. *Journal of Memory and Language, 35,* 606–621. http://dx.doi.org/10.1006/jmla.1996.0032

Slobin, D. I. (1973). Cognitive prerequisites for the development of grammar. In C. A. Ferguson & D. I. Slobin (Eds.), *Studies of child language development* (pp. 175–208). New York, NY: Holt, Rinehart & Winson.

St. Clair, M. C., Monaghan, P., & Ramscar, M. (2009). Relationships between language structure and language learning: The suffixing preference and grammatical categorization. *Cognitive Science, 33,* 1317–1329. http://dx.doi.org/10.1111/j.1551-6709.2009.01065.x

Thal, D. J., Bates, E., Zappia, M. J., & Oroz, M. (1996). Ties between lexical and grammatical development: Evidence from early-talkers. *Journal of Child Language, 23,* 349–368. http://dx.doi.org/10.1017/S0305000900008837

Thiessen, E. D., & Saffran, J. R. (2003). When cues collide: Use of stress and statistical cues to word boundaries by 7- to 9-month-old infants. *Developmental Psychology, 39,* 706–716. http://dx.doi.org/10.1037/0012-1649.39.4.706

Thompson, S. P., & Newport, E. L. (2007). Statistical learning of syntax: The role of transitional probability. *Language Learning and Development, 3,* 1–42. http://dx.doi.org/10.1080/15475440709336999

Toro, J. M., Pons, F., Bion, R. A. H., & Sebastián-Gallés, N. R. (2011). The contribution of language-specific knowledge in the selection of statistically-coherent word candidates. *Journal of Memory and Language, 64,* 171–180. http://dx.doi.org/10.1016/j.jml.2010.11.005

Weinert, R. (1995). The role of formulaic language in second language acquisition: A review. *Applied Linguistics, 16,* 180–205.

White, L., Valenzuela, E., Kozlowska-MacGregor, M., & Leung, Y. I. (2004). Gender and number agreement in nonnative Spanish. *Applied Psycholinguistics, 25,* 105–133. http://dx.doi.org/10.1017/S0142716404001067

Wilson, S. T., & Hudson Kam, C. L. (2013). *Learning syntax through statistics: When do transitional probabilities need a boost?* Atlanta, GA: Georgia Institute of Technology.

Yoshida, K. A., Pons, F., Maye, J., & Werker, J. F. (2010). Distributional phonetic learning at 10 months of age. *Infancy, 15,* 420–433. http://dx.doi.org/10.1111/j.1532-7078.2009.00024.x

(*Appendix follows*)

Appendix

SoftVoice Code for Words in Both Languages (Experimental and Control)

A	B	Prefix	Infix
Experimental language word type			
CHPAHZIYN	HHAAN	JUH/JUHB	MUWJ
TFOHBUWL	CHIYD	WOY/WOYS	NUWL
THMAXREY	GIYL		RIYG
VTIYSAA	NAHG		NEYD
AHBKAAG			
IYFJIYL			
AEVNAE			
OHNLUW			
Control language word type			
THRAHZIYN	HHAAN	JUH/JUHCH	MUWJ
ZWOHBUWL	CHIYD	WOY/WOYT	NUWL
KRAXREY	GIYL		RIYG
BLIYSAA	NAHG		NEYD
AHBKAAG			
IYFJIYL			
AEVNAE			
OHNLUW			

Received January 8, 2014
Revision received December 20, 2014
Accepted January 5, 2015 ■