
Volume 108, Number 2, August 2008
ISSN 0010-0277
pp. 303–600
ELSEVIER
COGNITION

Available online at www.sciencedirect.com
ScienceDirect
ELSEVIER
Cognition 108 (2008) 477–499
COGNITION
www.elsevier.com/locate/COGNIT

The curse of knowledge: First language knowledge impairs adult learners' use of novel statistics for word segmentation

Amy S. Finn *, Carla L. Hudson Kam

Department of Psychology, University of California, Berkeley, 3210 Tolman Hall #1650, Berkeley, CA 94720-1650, USA

Received 5 October 2006; revised 23 March 2008; accepted 15 April 2008

Abstract

We investigated whether adult learners' knowledge of phonotactic restrictions on word forms from their first language impacts their ability to use statistical information to segment words in a novel language. Adults were exposed to a speech stream where English phonotactics and phoneme co-occurrence information conflicted. A control where these did not conflict was also run. Participants chose between words defined by novel statistics and words that are phonotactically possible in English, but had much lower phoneme contingencies. Control participants selected words defined by statistics while experimental participants did not. This result held up with increases in exposure and when segmentation was aided by telling participants a word prior to exposure. It was not the case that participants simply preferred English-sounding words, however, when the stimuli contained very short pauses, participants were able to learn the novel words despite the fact that they violated English phonotactics. Results suggest that prior linguistic knowledge can interfere with learners' abilities to segment words from running speech using purely statistical cues at initial exposure.

Keywords: Language acquisition; Statistical learning; Learning constrains; Phonotactics; Transitional probabilities; Second language learning

* Corresponding author. Tel.: +1 510 642 7401.
E-mail address: amyfinn@berkeley.edu (A.S. Finn).

0010-0277/$ - see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.cognition.2008.04.002

1. Introduction

A great deal of recent research has been directed towards understanding statistical learning mechanisms. The picture that emerges is one of a very powerful, relatively domain independent, learning device that is deployed by both infants and adults (Hauser, Newport, & Aslin, 2001; Kirkham, Slemmer, & Johnson, 2002; Saffran, Aslin, & Newport, 1996; Saffran, Johnson, Aslin, & Newport, 1999; cf. Perruchet & Pacton, 2006). Still, there appear to be constraints on statistical learning, and some focus is being directed toward understanding what circumstances might mediate or constrain such learning (Bonatti, Peña, Nespor, & Mehler, 2005; Conway & Christiansen, 2005; Saffran, 2002). In the present paper we examine whether adults' knowledge of their first language might affect their ability to do statistical learning, in particular, whether first language phonotactic patterns affect the ability to segment according to consistent statistical regularities present in novel language stimuli when the two pieces of information about word boundaries conflict.

1.1. Background

Work in statistical learning has shown that adults and infants can track statistical regularities in running speech and use these regularities to find units akin to words, where a word is defined as a sequence of syllables that occur together consistently (e.g., Hauser et al., 2001; Saffran et al., 1996; Saffran, Newport, & Aslin, 1996; Toro & Trobalón, 2005). Learners in these studies are presented with a stream of syllables one after another. Some syllables occur together very consistently (for example, 100% of the time that syllable A occurs, syllable B follows it) and others less consistently (for example, syllable D only follows syllable C 33% of the time). Numerous studies have shown that learners are able to extract the words based on this statistical information after very little exposure to the stimuli (i.e., when there is no other information about words or word boundaries, save the contingencies). The statistical information used by learners in such studies is commonly assumed to be transitional probabilities (TPs), computed as the frequency of a pair of items (usually a pair of syllables) over the frequency of the first item in the pair.[1] In designs where frequency is not accounted for, however, the strings with high internal cohesion (high TPs) are usually also the most frequent strings, and thus it is possible that learners could be using either TPs or the simpler statistic co-occurrence frequency (string frequency) (see Perruchet, Tyler, Galland, & Peereman, 2004; Thompson & Newport, 2007, for discussions regarding the nature of these statistics). Regardless, demonstrations of statistical learning are robust and provide a possible explanation for how children might learn various aspects of language (see Seidenberg, MacDonald, & Saffran, 2002, for a more complete discussion).

[1] Recent computational work suggests that word segmentation cannot be achieved using transitional probabilities of syllables alone when the problem is scaled to a realistic setting (Yang, 2004). Work by Hockema (2006), however, suggests that this is less likely to be problematic when you look at phoneme transitional probabilities, which are strongly indicative of word boundaries in a realistic setting.

However, such learning is not unconstrained. It is now becoming clear that statistical learning is affected by a variety of factors. For example, Newport and Aslin (2004) found that learners were unable to learn dependencies between non-adjacent syllables, but could learn dependencies between non-adjacent consonants or vowels. Gómez (2002), however, found that learners could extract non-adjacent dependencies when there were pauses in between the items being tracked. Additionally, the relationships between the non-adjacent items were only acquired when there was a fairly large number of intervening items (Gómez, 2002). Saffran (2002) found a learning advantage for predictive dependencies in simultaneously, but not sequentially, presented visual stimuli. This was so even though learners were able to extract these same dependencies with sequential auditory presentation, suggesting that learning differs by modality (see also Conway & Christiansen, 2005). Attention can likewise affect learning outcomes. Toro, Sinnett, and Soto-Faraco (2005) showed that when learners' attention is directed away from the stimuli towards a distracter task that is either visual or auditory, learners fail to acquire the statistics present in the input stimuli (see also Turk-Browne, Jungé, & Scholl, 2005). Moreover, the level of difficulty of the distracter task affects how much learning is interfered with (Toro et al., 2005). It appears, then, that constraints can come from how the learning mechanisms interact with aspects of modality-specific perceptual processing or information processing more generally.

Another factor interacting with statistical learning, and the one we investigate here, is prior knowledge. Studies investigating word segmentation in infants have shown that older infants may ignore TPs when they conflict with learned probabilistic cues to word boundaries that exist in the infant's native language. For example, Thiessen and Saffran (2003) found that, when forced to choose, 6-month-olds will use TPs for word segmentation while 9-month-olds privilege the stress pattern that is characteristic of English bisyllabic words (strong–weak). Infants were exposed to a speech stream in which the syllable pairs with high between-syllable transitional probabilities exhibited a weak–strong stress pattern (the opposite of the preferred English pattern). The younger infants showed a pattern consistent with previous experiments; they segmented out the syllable pairs with high transitional probabilities. The older infants showed a quite different pattern. They segmented out items with lower syllable TPs but with the dominant English word-internal stress pattern. In a related study, Johnson and Jusczyk (2001) found that by 8 months of age, infants weigh co-articulation information more heavily than TPs.

What is surprising in these studies is not that the infants were able to use knowledge about word forms in their language to segment out novel words, but rather that they do so despite the presence of completely consistent information about syllable contingencies, something we know they willingly use at an earlier age. It is almost as if they stop performing statistical analyses. While this is likely beneficial to the young infant acquiring their native language, it can be thought of as acquired knowledge eventually interfering with statistical learning, at least when the two types of information conflict.[2] In

[2] Note that segmenting via stress cues will not always lead to extracting words, since some bisyllabic words in English have the opposite weak–strong pattern (especially verbs, Kelly & Bock, 1988). Interestingly, Jusczyk, Houston, and Newsome (1999) found that 7.5-month-old, but not 10.5-month-old, infants will sometimes mis-segment real speech when target words have the weak–strong pattern.

the present study, we investigate whether the same is also true for adults who are learning a new language. Do adults start learning a new system by initially deploying statistical learning, as do young infants when acquiring their native language? Or does the knowledge of their native language system interfere with statistical learning when confronted with a new language?

We examined this by exposing adult learners to an artificial language in which we pitted English word-formation patterns against TPs, and asked if they would segment words according to their knowledge of English or the TPs. In particular, the words, as defined by TPs, began with consonant clusters that violated English phonotactic rules. Phonotactics are the constraints a language places on the ordering of segments within words and syllables as well as information about distributions of things such as pronunciation variants. Like TPs, phonotactics are a type of co-occurrence information – co-occurrence given a particular position within a word, and thus, both likely arise from the same kind of learning (i.e., distributional learning). However, as with the stress cues mentioned previously, phonotactics are generalized from particular instances. Their use also involves different processes. Segmenting using TP information necessarily requires the learner to be tracking and computing things like frequency and co-occurrence. Segmenting via phonotactics, in contrast, involves the application of previously learned (i.e., previously computed) word-form regularities to novel stimuli. That is, the learner can, in some sense, segment out words by recognizing instances of patterns that are correlated with word boundaries in their known lexicon; it involves riding on the coat-tails of previously performed computations.[3] To be clear, although TPs and phonotactics are both types of co-occurrence information, they differ for our purposes in that TPs are statistical regularities specific to a particular speech stream, in this case the stimuli to which the participant is exposed during the experiment, whereas phonotactics are statistical regularities that are generalizations over speech to which a learner has been exposed, in this case the English to which the participant has been exposed over the course of her lifetime. Thus, TPs are a short-lived regularity, leading to segmentation but not necessarily stored beyond that process, in contrast to phonotactics which represent longer-term knowledge. Here our question is directly about whether this long term information can interfere with using TPs for word segmentation.

The extant literature makes conflicting predictions about what we should expect in our study. On the one hand are studies of interference in second language (L2) learners and users, which suggest that first language (L1) knowledge should interfere with the ability to segment novel words as defined by TPs (Marian & Spivey, 2003; Odlin, 1989; Weber & Cutler, 2006). On the other hand are studies showing: (1) that adults can learn novel phonotactic patterns (Dell, Reed, Adams, & Meyer, 2000; Onishi, Chambers, & Fisher, 2002; Weber & Cutler, 2006), and (2) that adults can consciously inhibit or deploy prior knowledge (Dienes, Altmann, Kwan, & Goode,1995). Unlike the work on interference these studies would suggest that adults should be able to overcome or ignore their L1 phonotactics to learn novel

[3] This is not to say that learners will not continue to perform distributional analyses on their lexicon as new words are added.

words as defined by TPs. Here we briefly review some of the relevant literature on both sides.

Learners tackling novel phonetic, morphological and syntactic systems often experience L1 interference (Odlin, 1989). That is, aspects of the L1's grammar negatively affect the learning of the same or similar features or structures in the L2. Although interference with respect to L2 acquisition is typically thought of mostly as a process affecting learning (and so something that, with time, can be overcome), it is a much more pervasive phenomenon. Studies have shown evidence for interference at the level of lexical access in highly proficient late bilinguals, for instance. In a study examining lexical competition in Russian–English bilinguals (L1 = Russian, L2 = English), Marian and Spivey (2003) found that lexical items from the L1 that shared word-initial phonological material with the L2 target word were activated and competed with the target for selection. This was the case even though the task was conducted completely in English. More closely related to our own study, Weber and Cutler (2006) found that highly proficient L2 speakers of English are still sensitive to the phonotactic patterns of their native language when listening to their L2. They used a word-spotting paradigm, in which participants listen to strings of nonsense in which real words are embedded and report whenever they detect a real word. The phonetic environment surrounding the words can be manipulated to make them easier or harder to spot, for instance, by embedding them such that the adjacent sounds cannot be in the same word according to the phonotactics of the language. In this study, Dutch–English bilinguals were asked to find embedded English (L2) words. Weber and Cutler found that participants were faster at detecting the English words when the L1 phonotactics were consistent with a word boundary precisely at the onset of the L2 target word. That is, the L1 phonotactics were leading them to posit a word boundary even though they were looking for words in the L2, a language with different phonotactic patterns. Both studies strongly argue for direct interference of first language knowledge even with proficient knowledge of L2. Based on the evidence just reviewed (as well as the infant studies reviewed earlier), then, we would expect at least some interference from L1 phonotactic patterns on word segmentation via statistical learning in a novel language.

Interference need not prohibit learning, however, and there is evidence that adults can learn novel phonotactic patterns. Although the speakers tested by Weber and Cutler showed interference from the L1 when performing a task in their L2, they were also faster at detecting the target word when the *English* phonotactics were suggestive of a boundary, indicating that they had learned the L2 phonotactic patterns. Moreover, there are studies showing that adults can rapidly learn new phonotactic patterns in the lab (Dell et al., 2000; Onishi et al., 2002). Onishi et al. (2002) had participants listen to C_1VC_2 'words' constructed such that individual consonants appeared only in either the C_1 or C_2 position, a type of phonotactic pattern. Training words were presented one at a time and participants were asked to repeat them. They were then tested on novel nonsense words in which the sounds either obeyed the same positional restrictions or violated them. Participants were faster when saying test words consistent with the training set than when saying test words which did not follow the same patterns, demonstrating that they had very quickly learned the novel positional restrictions. Although these studies say little about whether adults can extract words from running speech that violate their L1 knowledge upon initial exposure to the language (the question asked in the present work), they do demonstrate an important requisite ability: the ability to learn and maintain representations of word forms with phonotactics different from, and even sometimes in violation of, L1 patterns.

Another relevant factor is the fact that, unlike infants, adults can be told that they are leaning a novel language, and thus, might be able to direct their own learning. Gebhart, Aslin, and Newport (submitted for publication) presented adults with two different artificial languages one after another. When there was no explicit indication of the onset of the second artificial language, adults failed to segment the second language, but were able to segment the first. However, when explicitly told there were two languages and the switch in presentation was indicated with a pause, they were able to segment the second language. That is, a change in structure that participants did not notice and thus did not learn, was learned when they were made aware of its existence.

Adults also seem to be able to consciously apply implicit knowledge, suggesting that adults could chose to suppress their L1 system at learning. Dienes et al. (1995) conducted an Artificial Grammar Learning experiment in which they sequentially exposed participants to strings generated by two different grammars. That is, participants first learned the strings generated by grammar one, then those generated by grammar two. At test, participants were asked to respond only to novel grammatical strings from one of the grammars. Participants were able to do this despite having no explicit knowledge of either the target (to be selected) or non-target (to be ignored or suppressed) grammar. There is also a great deal of work showing that bilinguals can actively inhibit or suppress their knowledge of one or the other of their languages (also implicit knowledge), although it is not clear how conscious this process is (Meuter & Allport, 1999; Rodriguez-Fornells, Rotte, Heinze, Noesselt, & Münte, 2002; Rodriguez-Fornells et al., 2005).

But such selection and suppression need not be so directed. As mentioned earlier, Marian and Spivey (2003) found that when participants were unaware that their bilingualism was relevant to the experiment, there was interference from the L1 into the L2 lexicon, but not the L2 into the L1 lexicon. However, in another study (Spivey & Marian, 1999) they found that when participants were aware that they were involved in an experiment specifically about bilinguals, the pattern of results was reversed: there was interference from the L2 into the L1 lexicon, but not from the L1 into the L2. Thus, the degree of activation of L1 and L2 knowledge (and thus interference) appears to be dependent on situational and contextual variables that can be manipulated.

Together, these studies suggest that adults learning a new language could withhold knowledge of their native language, especially when explicitly made aware that they are learning something novel. Thus, unlike infants who cannot be told they are learning a 'new' language, and so go with the learned cues over TPs in laboratory tasks, adults can be made aware of their task and this in turn, may allow them to selectively suppress their L1 knowledge so that statistical learning abilities are free to operate.[4] The present study asks if adult learners are indeed able to withhold L1 knowledge and begin to learn a new system by tracking TPs. If adult learners use TPs and ignore phonotactic regularities from their L1, this would suggest that they can selectively inhibit L1 knowledge (or that L1 knowledge is simply not a factor affecting statistical learning in adults, a possibility that the interference literature would suggest is unlikely). If instead they parse the speech stream according to English phonotactics, it would indicate that their prior knowledge is interfering with statistical learning, and constraining learning in ways not beneficial for successful acquisition.

2. Experiment 1

2.1. Method

2.1.1. Participants

Forty undergraduates at the University of California, Berkeley participated for course credit. Twenty were in the experimental condition, 20 were in the control condition. In this and all other experiments, participants were native speakers of English who reported normal hearing.[5]

2.1.2. Stimuli

Experimental and control stimuli both consisted of eight two syllable words (CCVCV), each beginning with a consonant cluster. For the experimental stimuli, these CC onsets violate the word-initial phonotactic rules of English. In the control stimuli, CC onsets are licit.

The experimental stimuli were: /tfobu/, /čpʌzi/, /btegʌ/, /kmodu/, /vtisa/, /fselo/, /psune/, and /θmʌre/. Note that all but two of the experimental clusters (čp, vt) can occur as word internal strings. The eight control words were: /zwobu/, /θrʌzi/, /kregʌ/, /plodu/, /blisa/, /vrelo/, /twune/, and /stʌre/.[6]

Each word was generated with the text-to-speech program SoftVoice (Katz, 2005). The synthesizer produced syllables with a monotonic F_0 (fundamental frequency) of 83.62 Hz. All vowels were matched for length and there were no co-articulation effects. We used synthesized speech to allow better control of the above-mentioned parameters. Use of natural speech risks the inclusion of additional

[4] Interestingly, there is one study in which infants were in essence trained to overcome their learned stress pattern preferences (Thiessen & Saffran, 2007).

[5] Participants were considered native speakers of English if they were exposed to the language from birth and began speaking English before the age of four. Bilinguals were allowed. As a check, we collected language background information from all participants, and examined grammars of all languages other than English to see whether they allowed any of the clusters we used. French was the only language spoken by any of our participants ($n = 2$, both late learners) that allowed onsets present in our experimental stimuli. Post-hoc analyses showed that their performance did not differ from the other participants, even on the items including the French clusters.

[6] Two of these are not common, primarily occurring in loanwords (e.g. **zw**eiback) and sound effects (e.g. **vr**oom).

segmentation cues through varying degrees of co-articulation, different vowel lengths, amplitudes, frequencies, etc.

Words were presented quasi-randomly with no pauses and no immediate repetitions, yielding transitional probabilities of 1.0 for syllable transitions that are word internal and .143 at word boundaries. Importantly, phoneme transitional probabilities (PTPs) within words were higher than those across word boundaries, with word-internal PTPs, ranging from .25 to 1.0 and PTPs across word boundaries ranging from .035 to .143. Thus, there was no overlap in the within-word and across-word PTPs. (The ranges are a consequence of using some phones multiple times, in an attempt to create stimuli that is somewhat more naturalistic than that which is typically used. See Saffran et al., 1996, for previous research comparing ranges of TPs.) During exposure, each word occurred 560 times, with exposure lasting a total of 17 min 59 s. A sample stretch of input is presented here: /tfobuc̆pʌzibtegʌkmoduvtisafselopsuneθmʌre/.

2.1.3. Tests

After exposure participants were given a forced-choice test. Test items were of two types: (1) word vs. non-word, and (2) word vs. split-cluster word. In both tests, words were the words (with 1.0 word-internal syllable TPs and high PTPs) to which participants had been exposed. Non-words consisted of the first syllable from one word paired with the second syllable from a completely different word, e.g., /kmone/. Although participants had heard each of the two syllables in the non-word an equal number of times, they had never heard the two syllables in succession. These test items were included to serve as a control, indicating whether participants were attending to the stimuli and at least tracking probabilities at the level of the syllable. In principle, these items tell us little about participants' actual segmentation – that is, where they think the word boundaries are – since both correct and incorrect options contained the violating clusters, and so we expect both experimental and control groups to perform equally well on these test items. (Of course, it is possible that the clusters cause so much difficulty for the experimental participants that they fail to learn even the syllabic TPs, which would predict that they should select randomly on these items, unlike participants in the control group.)

The second test item type more directly assessed participants' word segmentation. In these items participants were asked to compare a word with an exposure word minus the first consonant with another word's initial consonant at the end (resulting in a CVCVC structure, e.g., /pʌziθ/). In essence, we constructed the split-cluster words by shifting one phoneme to the right in the exposure stimuli, making these much like 'part-words' (that are shifted an entire syllable to the right for tri-syllabic words) as in other statistical learning designs (see e.g., Saffran et al., 1996). We refer to these as split-cluster words because we split the consonant cluster at the beginning of the word, the result being a viable English word. While the non-word test items are designed to probe learners' knowledge of transitional probability at the level of the syllable, these items probe knowledge of transitional probability at the level of the phoneme. The split-cluster words have lower PTPs than the words, but are licit according to the rules of English. For example, although /pʌziθ/ does not violate

English phonotactics, the θ (th) sound only follows the phone /i/ 7.14% of the time in the exposure stimuli. Thus, we created test items that should sound better according to participants' prior knowledge of English, but which are not words according to the PTPs in the new language.[7]

There were 8 of each type of test item (word vs. non-word, word vs. split-cluster word), yielding 16 test items in total. Test items of the two types were interleaved quasi-randomly and two versions of the test were created. (One version was the reverse of the other.) Test stimuli were generated with exactly the same procedure as exposure stimuli. The two items in a pair were presented one after another with a 1-s pause in between. There was a 3-s pause in between pairs during which participants were expected to answer.

2.1.4. Procedure

Participants were told that we were testing differences between adults' and children's abilities to learn languages. They were told that they were going to listen to a new language, and that after exposure they would be tested to see what they had learned about the language. They were instructed simply to listen to the speech as best they could, but not to over-think or ignore what they were listening to. To encourage this, participants colored using crayons or markers during exposure. Participants wore noise-cancelling headphones during presentation and testing.

After exposure, participants were given the forced-choice test. They were told that they would listen to pairs of possible words and were asked to "choose which is a better example of a word in the language" they were exposed to. They were encouraged to make their best guess if unsure. Participants indicated responses on an answer sheet, circling 1 if the first item in the pair sounded better and 2 if the second one sounded better. After the test, participants completed a survey probing their demographic and language backgrounds.

2.2. Results and discussion

Fig. 1 shows participants' performance on non-word (black bars) and split-cluster test items (white bars). In this and all other figures, error bars represent standard error. Data are shown separately for experimental and control participants. Performance on the non-word items was good; participants in both conditions chose words over non-words more often than chance (one sample t-tests – experimental participants: $t(19) = 9.19$, $p < .001$; control participants: $t(19) = 4.46$, $p < .001$). Moreover, the two groups did not differ from each other on this measure (independent samples t-test: $t(38) = 1.25$, $p = .22$), indicating that they were generally able to track statistical information at least at the level of the syllable in both sets of stimuli. (In this and all analyses reported in this paper, chance performance is 50%.)

However, it appears that participants exposed to onsets that violate English phonotactics did not correctly segment the stimuli, that is, they did not extract words defined

[7] Note that this additional sound should make the task easier, since the PTP between the second vowel and the added consonant is always quite low. Thus, it should provide a very strong cue to the participant that the split-cluster word is not the correct response.

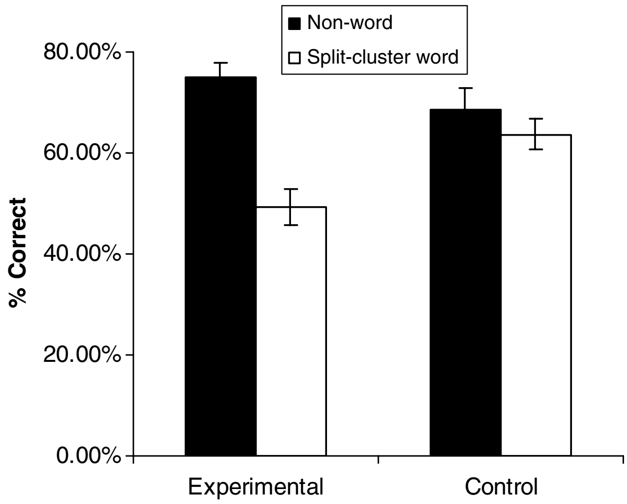

Fig. 1. Experiment 1: Percent correct on word vs. non-word and word vs. split-cluster word items by condition.

by PTPs; a t-test on the split-cluster vs. word test items showed that participants in the experimental condition did not choose words more often than chance ($t(19) = -.175$, $p = .863$). Their performance cannot be explained by a more general problem tracking transitional probabilities over phonemes or with consonant clusters as onsets, however, since participants in the control condition chose words over split-cluster words more often than chance ($t(19) = 4.40$, $p < .001$). Nor is it the case that experimental participants are simply choosing the test items that accord with English phonotactic rules (the split-cluster word), which would result in performance significantly below chance. As Fig. 1 makes apparent, experimental and control groups' performance on this measure was significantly different ($t(38) = -3.03$, $p = .004$).

To summarize, performance on non-word test items shows that adults can learn transitional probabilities at the level of the syllable in a novel language that violates the phonotactic constraints of their native language. However, this test does not indicate whether participants learned the PTPs, that is, whether they extracted correctly segmented words, clusters and all. This was addressed by the split-cluster items. Performance on these items was at chance for the experimental group, but well above chance for the control group. Overall then, these results suggest that prior linguistic experience seems to interfere with participants' ability to correctly segment the words, in this way constraining the operation of statistical learning. In three additional experiments we examine whether this effect of experience can be overcome.

3. Experiment 2

In Experiment 1 learners were only exposed to speech stimuli for about 18 min. It could be that with increased exposure the effect of L1 phonotactics is reduced. We therefore extended exposure to 72 min, quadrupling exposure, and split the exposure period into two sessions occurring over 2 days.[8] Evidence suggests that memory is enhanced in adults and infants by sleep (Gómez, Bootzin, & Nadel, 2006; Plihal & Born, 1997; Stickgold & Walker, 2005). It could be that with increased exposure and time for sleep (and perhaps memory consolidation, see e.g., Jenkins & Dallenbach, 1924; Maquet, Peigneux, Laureys, & Smith, 2002), the constraining effect of prior knowledge is reduced.

3.1. Method

3.1.1. Participants

Forty undergraduates at the University of California, Berkeley participated for course credit. Twenty were in the experimental condition, 20 were in the control condition.

3.1.2. Stimuli and tests

The stimuli and test items for this experiment were the same as Experiment 1, except that the total exposure was quadrupled, totaling 72 min.

3.1.3. Procedure

Save for the extension of exposure time, the procedure was exactly the same. Day one consisted of 36 min of exposure. Day two was an additional 36 min followed by the test. Sessions were always completed on consecutive days.

3.2. Results and discussion

The results for non-word and split-cluster test items for both control and experimental subjects are shown in Fig. 2. Participants exposed to the control stimuli consistently chose words over non-words ($t(19) = 7.73$, $p < .001$) and split-cluster words ($t(19) = 6.68$, $p < .001$). Participants in the experimental condition, however, showed a different pattern. A t-test showed that participants choose words over non-words more often than chance ($t(19) = 9.45$, $p < .001$). However, a t-test on the split-cluster vs. word test items showed that, yet again, participants did not choose words more often than chance ($t(19) = 1.60$, $p = .126$). As in Experiment 1, control and experimental participants' performance did not differ for non-word items ($t(38) = -.67$, $p = .507$) but did for split-cluster items ($t(38) = -2.84$, $p = .007$). The pattern of results from Experiment 1 was not altered even when exposure time was quadrupled and subjects were given an opportunity to sleep between exposure periods.

[8] With a separate group of 20 participants, we first tried extending exposure to 36 min (double Experiment 1) and found no improvement: non-word test items were significantly different from chance ($t(19) = 2.38$, $p = .028$), but split-cluster items were not ($t(19) = .65$, $p = .522$).

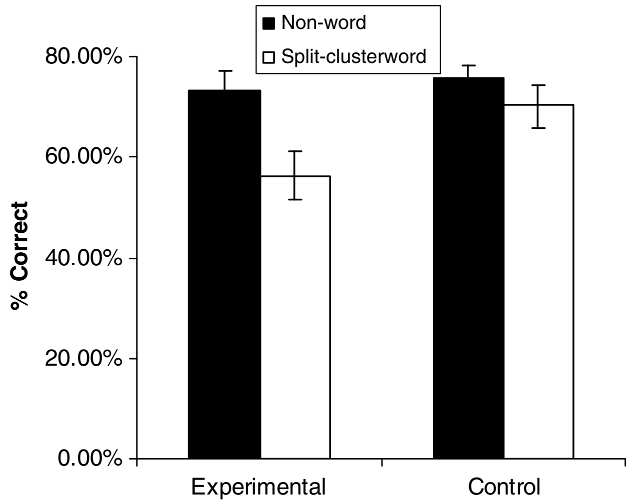

Fig. 2. Experiment 2: Percent correct on word vs. non-word and word vs. split-cluster word items by condition.

4. Experiment 3

The previous experiments show that the constraining effect of prior knowledge is quite robust. This could, however, be a product of this artificial exposure situation. In reality, it is unlikely that learners are exposed to completely un-parsed stimuli for so long. Learners will generally hear at least a few words in isolation (Brent & Siskind, 2001). It has been argued that these initial words are what learners use to make initial generalizations about the prosody and phonotactics of their language – generalizations they later employ to further parse the speech stream (see e.g., Werker & Yeung, 2005). It has also been argued that having a few initial words can provide anchors in the speech stream, trimming it down and generally easing the task of segmentation (Bortfeld, Morgan, Golinkoff, & Rathbun, 2005).

To investigate whether this would help adult learners overcome the constraining effect of their existing linguistic knowledge, we gave participants one of the words before exposure to see if this would improve their segmentation performance. By giving participants explicit information about one of the words, we are providing them with clear information about the boundaries of other words as well; if you know where one word ends, you also know where another begins.

4.1. Method

4.1.1. Participants

Twenty undergraduates at the University of California, Berkeley participated for course credit.

4.1.2. Stimuli and tests

The stimuli and test items for this experiment were the same as Experiments 1 and 2.

4.1.3. Procedure

The procedure for Experiment 3 was exactly the same as for Experiment 1, except that after the more general instructions participants were told "/kmodu/ is a word in the language you are about to listen to." As the experimenter was about to leave the room and just before exposure began, she said "remember /kmodu/ is a word in the language." (/kmodu/ was chosen because it was easy for experimenters to pronounce.)

4.2. Results and discussion

Fig. 3 shows participants' performance on non-word test items, split-cluster test items other than /kmodu/ (the word we gave participants prior to exposure), and /kmodu/ itself. Somewhat surprisingly, the results were consistent with the previous experiments. Participants chose words over non-words more often than chance ($t(19) = 4.19$, $p < .001$), but did not select words more often than split-cluster words when /kmodu/ was not considered ($t(19) = 1.10$, $p = .286$). Performance on that item was at ceiling; all participants chose /kmodu/ with 100% accuracy. It appears that even when participants are given a full word before exposure, it does not dampen the effect of prior knowledge when segmenting novel stimuli. This is true even though participants remembered the word we gave them before exposure and could have used this to inform segmentation for other items.

5. Experiment 4

One possible explanation for the pattern of data reported in Experiments 1–3 is that participants sometimes chose split-cluster words at test because they were basing their decisions on English, not because they failed out correctly segment out the words during exposure. That is, the data reflects something happening at the time of testing, not what they have or have not learned about the experimental stimuli. It is worth pointing out that this is unlikely to explain the data in full because participants were at chance on split-cluster words; if decisions were based completely on

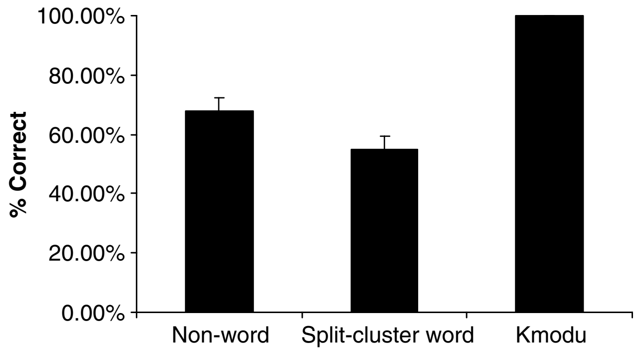

Fig. 3. Experiment 4: Percent correct on word vs. non-word items, word vs. split-cluster items, and kmodu.

English, they would be consistently below chance on this test. Still, knowledge of English could have played a disproportionately stronger role during test/decision than exposure. To rule this out as a possible explanation for the previous data, we conducted one further experiment. We exposed participants to the same experimental and control stimuli with short pauses inserted between words, in essence, providing other simpler cues to segmentation. If knowledge of English has an impact on segmentation during exposure, and not just during the test/decision phase, the performance of the control and experimental groups should be equal and above chance. If however, knowledge of English interferes selectively at test, then the pattern of results reported in Experiments 1–3 should be replicated, since the same conditions at test apply.

5.1. Method

5.1.1. Participants

Thirty undergraduates at the University of California, Berkeley participated for course credit – 15 each in the experimental and control conditions.

5.1.2. Stimuli and tests

Control and experimental stimuli and test items for this experiment were the same as for Experiment 1 except that during exposure a short silence (400 ms) was inserted between each word. Test items were exactly the same.

5.1.3. Procedure

The procedure for Experiment 4 was exactly the same as for Experiment 1, except that participants completed the test on a computer and indicated their choice with a button press.

5.2. Results and discussion

Fig. 4 shows participants' performance on non-word (black bars) and split-cluster test items (white bars). Data are shown separately for experimental and control subjects. Performance on the non-word items was very good; participants in both conditions choose words over non-words more often than chance (one sample *t*-test for experimental participants: $t(14) = 10.22$, $p < .001$; and control participants: $t(14) = 6.97$, $p < .001$), indicating that they were generally able to track statistical information in both sets of stimuli. Similarly, performance was also very good for the split-cluster vs. word items (experimental participants: $t(14) = 6.35$, $p < .001$; and control participants: $t(14) = 12.25$, $p < .001$). Importantly, independent samples *t*-tests show that performance did not differ across control and experimental groups for either the split-cluster ($t(28) = -1.05$, $p = .301$) or the non-word test items ($t(28) = 1.26$, $p = .217$). Taken together with results from Experiments 1–3, this indicates that participants' knowledge of English does not have an impact selectively during the decision process at test. Rather, knowledge of English is likely to have a more global impact during exposure. Additionally, above chance performance

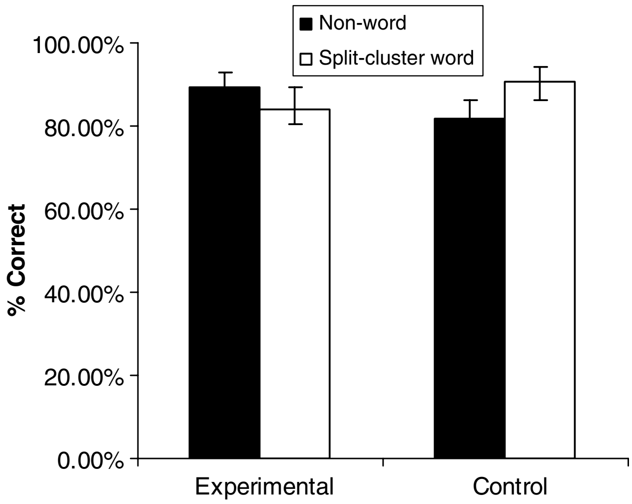

Fig. 4. Experiment 5: Percent correct on word vs. non-word and word vs. split-cluster word items by condition.

on both of these tests suggests to us that learners are indeed able to perceive clusters in both experimental and control stimuli.[9]

6. General discussion

In this paper we investigate the possibility that prior knowledge can mediate and interfere with statistical learning. In particular, we examined the impact of learners' knowledge of L1 phonotactics on statistical word segmentation. In three experiments we exposed participants to stimuli containing words defined by transitional probabilities but that started with consonant clusters which are not possible word-initially in English, the native language of all our participants. We found that participants did not extract the words when their input contained these illegal word-initial clusters. This result held despite substantial increases in exposure, and even when we explic-

[9] To further investigate the perceptability of the experimental clusters, we had 10 additional people participate in a perception test. They heard either one of the experimental words or one of eight fillers, and were asked to circle one of two possible orthographic representations corresponding to what they heard. The fillers were all monosyllabic, but varied in their syllable structure so as not to highlight the nature of the test. In each case, participants' choice was between items with and without a cluster (e.g., kmodu or modu, ren or rens). Filler items sometimes had initial or final clusters and sometimes had no clusters. This was to alleviate the concern that subjects would simply learn to choose orthographic representations with clusters. Somewhat counter-intuitively, three of the filler items contained clusters that violated the phonotactic patterns of English. We did this to avoid the test containing different patterns in the mono- and bi-syllabic items. Performance was good for both experimental (81.3% correct) and filler (95% correct) items. Both are significantly above chance (experimental: $t(9) = 5.24$, $p < .001$; filler: $t(9) = 22.05$, $p < .001$). Performance on the filler items was significantly better than performance on the experimental words ($t(9) = -2.7$, $p < .05$). We believe better performance on filler items may be due in part to the fact that they are monosyllabic and therefore more easily recalled.

itly told participants one of the words prior to exposure, in essence, giving them some information about segmentation. In an additional experiment, we demonstrated that native speakers of English do not have trouble at test when the exposure stimuli are constructed with silences in between the words; this suggests both that learners can perceive these illegal word-initial clusters and that knowledge of English is having an impact during segmentation, rather than interfering with participants' ability to express their knowledge at test. This pattern of results suggests to us that prior linguistic knowledge can mediate and therefore constrain the operation of basic mechanisms of acquisition, in particular, whether statistical learning is employed for word segmentation.

Interference from a previously learned language is a very well established phenomenon (see Odlin, 1989, for a comprehensive review), and so at first glance, our findings are not that surprising. Typically, however, work on transfer and interference has focused on how aspects of an L1 grammar affect the learning of the same feature or structure in a second language, for example, whether the way relative clauses are formed in the L1 affects learners' acquisition of relative clause structures in the L2. This is not the nature of the interference we were investigating here – we were not looking at how knowledge of L1 phonotactics interferes with the ability to acquire phonotactic restrictions in a new language, at least not directly. We were looking at how L1 phonotactics can interfere with the learners' ability to acquire something else – word boundaries. An inability to find word boundaries could then affect a learners' ability to acquire new phonotactic restrictions, of course. Given that the discovery of word boundaries necessarily precedes the formation of rules concerning those boundaries, if one has difficulty finding them, learning phonotactic generalization about those boundaries will be difficult. Thus, L1 phonotactic knowledge may interfere with the acquisition of L2 phonotactics in an indirect, as well as a direct, way.

6.1. Alternative interpretations

There are a few other potential interpretations of our results that deserve some discussion. First, we used a more complicated word segmentation test than previous studies. In particular, the word vs. split-cluster word test asked participants to choose between items that differed only in terms of two phones, one at the beginning and another at the end. It could be that learners are simply not able to track statistics at the level of the phoneme. Newport and Aslin (2004), however, have shown that learners can track non-adjacent phonemes. Furthermore, our control subjects – who heard words containing word-initial clusters that are licit in English – did not have difficulty with this type of comparison. Unlike the experimental participants, the control participants correctly segmented the words from the speech. Another possibility is that many of our violating cluster onsets are simply linguistically unnatural, and thereby impossible due to the constraints of Universal Grammar. Seidl and Buckley (2005), however, have shown that infants can learn arbitrary patterns that violate linguistic universals when the test relies on perception rather than production (as ours does). Naturalness is therefore unlikely to be the reason our participants failed to learn the word boundaries. Also relevant is the fact that participants exposed to the speech stream containing pauses managed to learn the correct word forms, illicit (and possibly unnatural) clusters and all.

We should point out that although we have discussed statistical word segmentation in terms of transitional probabilities, given that the correct answers were more frequent than the wrong answers during exposure (i.e., /blisa/ vs. /lisav/), participants could have made their decisions at test based simply on how often they heard each string in its entirety together and not the difference in the PTP at the end of the word (see e.g., Kinder & Assmann, 2000; Meulemans & Van der Linden, 1997). This point pertains mostly to the control conditions, as experimental participants did not choose the more frequent string more often than chance (except when the input was pre-segmented for them, as in Experiment 4). This possibility, however, does not negate our main point, which is that L1 knowledge interferes with learners' abilities to acquire the statistics related to word boundaries in a novel language – whether these statistics are string frequency or TPs – since participants exposed to the experimental language did not manage to correctly segment the words.

Although we designed our study in terms of speakers' knowledge of possible and impossible word forms in their L1, our learners could have been using a different flavor of phonotactic knowledge for segmentation. English speakers do not only know that our experimental CC onsets are illicit in word-onset position according to English (our intended manipulation), but also that when they do occur in English, they are highly predictive of word boundaries. It is possible that it is this other, more probabilistic, aspect of their L1 knowledge that is at work. On this explanation, participants hear the experimental clusters and assume that a word boundary exists between the consonants because a boundary typically occurs there, not because they know that those two consonants cannot occur in onset position.[10] Similarly, they learn the control words because they expect the two consonants to occur in the same word, not because they have no prior knowledge that forces them to segment one way or the other. This is also phonotactic information, however, it is not the phonotactic information we were intending to test. While we did not control for this variable directly (it is quite difficult to equate this for licit and illicit clusters due to the facts of English), there is some variation on exactly this parameter in both the experimental and control clusters we used, allowing us to examine whether or not this played a role in participants' segmentation.

In particular, we examined whether participants had more trouble correctly segmenting items containing clusters that are highly predictive of word boundaries and an easier time with clusters that typically occur within words. To establish the likelihood of a word boundary occurring within each of the 16 diphones (i.e., clusters), we searched two data sets of American English from the CHILDES database:

[10] Using the latter type of information for word segmentation would seem to involve more of a Baysian-type computation along the lines of, 'given this sequence of sounds what it the likelihood that they are in one versus in two words'. However, we do not wish to imply that this makes them less likely given the emerging evidence that adults and infants can perform just these sorts of inferences (especially as there is at least one study showing that older infants will segment artificial speech streams along just these lines, Mattys, Jusczyk, Luce, & Morgan, 1999).

the Brown (1973) and Gleason, Perlmann, and Greif (1984) corpora. These corpora cover speech directed to children from 1 year 6 months of age to 5 years, 2 months, and include substantial amounts of adult-to-adult speech as well. We examined all speech not produced by the child. In all, the utterances examined included 545,637 words. We searched for every occurrence of the diphone, and coded each occurrence for the presence (or absence) of a word boundary. (Note that to be counted as a diphone, the two sounds had to occur in the same utterance.)

For the experimental clusters, the probability of a word boundary occurring within the diphone was quite high (/tf/ = .98, /čp/ = 1.0, /bt/ = .75, /km/ = .96, /vt/ = 1.0, /fs/ = .6, /ps/ = .11, /θm/ = .99). The opposite was true for the control clusters (/zw/ = .99, /θr/ = .008, /kr/ = .038, /pl/ = .018, /vr/ = .147, /tw/ = .94, /st/ = .087). However, as one can see, there are some exceptions to this pattern. For the experimental stimuli, /čp/ and /vt/ never occurred within a word and /tf/, /km/, and /θm/ did so very rarely; yet for two of these (/tf/ and /čp/), participants in all experiments segmented these items correctly at above chance rates. Moreover, one of the experimental clusters occurred fairly evenly in both contexts (/fs/) and another almost always occurred within a word (/ps/). If participants were segmenting based on the likelihood of a diphone occurring within a single word vs. two words in English, then words containing these two clusters should have been easily learned. But this was not the case. Likewise, our search showed that two of the control clusters were highly predictive of word boundaries (/zw/ and /tw/) and yet, control participants learned these words with ease. Thus, learners did not have more difficulty segmenting words beginning with clusters that are highly predictive of word boundaries in English. We therefore do not think that this is driving the pattern of results reported above.[11]

6.2. How do L2 learners find words?

Our data strongly suggest that L1 knowledge mediates the functioning of a mechanism of acquisition – statistical learning. This begs a very important question: how do adult L2 learners ever manage to learn word boundaries? There are several answers to this question. First, explicit learning (and teaching) may be very important, enabling learners to acquire words with phonological combinations not allowed in their native languages. Some evidence for this comes from Experiment 3, where participants had no difficulty with /kmodu/, the word they had been given explicitly.

[11] This of course raises the question of what is driving the variation in performance, that is, why are words beginning with /tf/ and /čp/ consistently segmented at above chance rates in all experiments. (These were the only words that showed any consistent patterns across all experiments.) As just explained, it is not correlated with the probability of a word boundary occurring in the diphone. Following Saffran, Newport et al. (1996) we checked to see whether words with higher average PTPs were segmented correctly more often than those with lower average PTPs. Again, the results were negative; this does not explain the variation. We also examined the data from the perception experiment (discussed in a previous footnote) to see whether there was a correlation between performance on segmentation and perceptability; there was not ($r = .304$). The two words that consistently led to above chance performance contained clusters that were correctly perceived by 90% of the participants. However, three other clusters that were also correctly perceived 90%+ of the time were in words that were not well learned by participants.

Additionally, when words are pre-segmented during exposure (as in Experiment 4), participants do learn them. This also might explain the apparent contrast between our results and those of Onishi et al. (2002) – they found that adults could learn novel phonotactic patterns very quickly in the lab whereas the adults in our study had so much trouble with words containing phonotactic patterns different from English; however, their participants heard pre-segmented words[12]. Second, not all words in any given L2 will violate L1 phonotactic patterns, and so the phenomenon we have shown will not be a factor all of the time; very often, learners will be able to use transitional probabilities to find word boundaries.

Increasing exposure may also lead learners to better learning of word boundaries. Work explicitly looking at exposure to multiple artificial languages, where the input language actually changes during exposure seems to support this notion. As mentioned previously, adult learners do not seem to notice the switch and learn the second language without explicit instruction and cues. Substantially increasing the exposure to the second (or switched) language does seem to lead to some learning, however (Gebhart et al., submitted for publication). (Although learners in these studies still demonstrate significantly more knowledge of the first language.) Similarly, our data hint to a trend in this direction, although the differences are not significant (performance on split-cluster items, Experiment 1 vs. Experiment 2, experimental participants ($t(38) = -1.3$, $p = .201$); control participants ($t(38) = -1.49$, $p = .143$). Our longest exposure condition, 72 min over 2 days, is still very little exposure as compared to natural language exposure. It is a great deal more than typically seen in statistical learning studies, however, and other studies with long exposures have shown that things that are not learned in short exposures are not necessarily learned with longer exposures either (e.g., Newport & Aslin, 2004). The phonotactic restrictions we explore likewise seem to be persistent in spite of large increases in exposure. Weber and Cutler's (2006) work, discussed previously, shows that L1 phonotactic restrictions affected participants' ability to find embedded L2 words even in highly competent L2 speakers, who have had a great deal of L2 exposure.

6.3. Prior knowledge and statistical learning

The constraining impact of L1 knowledge demonstrated thus far in no way undermines previous empirical demonstrations of adults' statistical computations over novel linguistic input. Clearly adults can do this. Indeed, positional restrictions on phonemes, the kind of phonotactics we have been exploring in the present work, involve statistical regularities, and we know that adult language learners can learn these (see e.g., Dell et al., 2000; Onishi et al., 2002; Weber & Cutler, 2006). Moreover, adults' computational abilities are not limited to learning phonotactics. Many studies, including our own, have shown that adults can perform the relevant compu-

[12] This is not the only difference that might be relevant. They were asking people to learn further restrictions, whereas in our study adults were being asked to essentially unlearn a restriction. That is, learning that something cannot go where you thought it could, as opposed to it can go where you thought it could not. (This is also true of Warker & Dell, 2006).

tations to do statistical word segmentation, at least at the level of the syllable (e.g., Saffran, Newport et al., 1996). The novel contribution of this work is to show that adults' knowledge of their native language system can interfere with statistical learning when such L1 knowledge conflicts with statistical cues to word boundaries. This is true even though they are aware that they are engaged in a learning task, unlike the infants in previous studies that pitted learned cues against TPs (Johnson & Jusczyk, 2001; Thiessen & Saffran, 2003). Because of previously acquired knowledge, the learner fails to see the optimal pattern in the input (the TPs). Importantly, the failure cannot be attributed to an inability to perform the computations in question; we know that both adults and children can segment based on TPs (Saffran et al., 1996; Saffron, Newport et al., 1996). Rather, it is an inability to see through the conflicting information (in this case phonotactic information) that leads to a lack of learning. We fully expected adults to go back to this initial learning stage by increasing exposure, giving them sleep, and even telling them a word before exposure. Future research is needed to uncover the conditions that might lead adults to override knowledge from their native language.

Recent theory from the L1 literature can help us understand why adults might show such strong interference from their L1 on statistical learning when they seem to otherwise be very good at tracking statistical information. According to Kuhl's Native Language Neural Commitment hypothesis (Kuhl, 2004), our brains make an early commitment to the statistical and prosodic regularities in language, and this commitment is necessary for learning higher order native-language information. This hypothesis explicitly predicts a greater degree of difficulty in processing a non-native language that does not adhere to these more basic regularities, something demonstrated in the present data. Phoneme and phonotactic learning is an essential component of these initial computations. In several landmark studies, Werker and Tees (1984a,b) showed that older infants loose the ability to discriminate many non-native phonemic distinctions, likely as a consequence of growing expertise with one's native language. Considering this and other empirical work on phonological acquisition showing early L1 tuning (Aslin, Pisoni, Hennessy, & Perey, 1981; Eimas and Miller,1992; Eimas, Siqueland, Jusczyk, & Vigorito, 1971; Kuhl, Williams, Lacerda, Stevens, & Lindblom, 1992; Werker & Lalonde, 1988), Werker and Tees (2005) have proposed a cascading model of language acquisition. On their model (as in Kuhl's), this early phonetic category tuning impacts and must precede other aspects of phonological learning (including phonotactic learning), and in turn many other aspects – if not all – of language acquisition. Our adults, therefore, may have particular difficulty using statistics to segment words in a novel language because the (conflicting) information from the L1 – phonotactics – is part of the crucial scaffolding that has led them to learn their L1.

These theories then begin to help us understand a conundrum: animals, adults, infants and the like are all capable of many computations, as experiment after experiment has demonstrated (Hauser et al., 2001; Kirkham et al., 2002; Saffran et al., 1996, 1999), so why is it that we have so easily shown a failure? Statistical learning, although a form of domain-independent learning, operates over information that can be very much *domain-specific*. Previous work has shown that statistical learning can be influenced by how we process information in particular domains (e.g., Conway & Christiansen, 2005; Saffran, 2002). Kuhl (2004) and Werker and Tees (2005) point out that how we process information in language is affected by what we already know, and we show that this too has implications for statistical learning.

Acknowledgments

This research was funded in part by a National Science Foundation GRFP Award to A.S.F. We thank O. Hemmy-Asamsama, J. Morrison, A. Klem, J. Newman, I. Lee, E. Cox, and N. Kramer for their help in conducting these experiments. Thanks also to E. Thiessen and the UC Berkeley Phonetics Lab for advice regarding stimuli construction. We are also grateful to A. Flevaris and T. Beyer for comments on earlier drafts of this paper. We are very grateful to three reviewers for their very thoughtful and helpful suggestions. We are particularly thankful to all of the members of the UC Berkeley Language and Learning Lab for their ongoing help and support during every stage of this project.

References

Aslin, R. N., Pisoni, D. B., Hennessy, B. L., & Perey, A. J. (1981). Discrimination of voice onset time by human infants: New findings and implications for the effects of early experience. *Child Development, 52*, 1135–1145.

Bonatti, L. L., Peña, M., Nespor, M., & Mehler, J. (2005). Linguistic constraints on statistical computations. *Psychological Science, 16*, 451–459.

Bortfeld, H., Morgan, J. L., Golinkoff, R. M., & Rathbun, K. (2005). Mommy and me: Familiar names help launch babies into speech-stream segmentation. *Psychological Science, 16*, 298–304.

Brent, M., & Siskind, J. M. (2001). The role of exposure to isolated words in early vocabulary development. *Cognition, 81*, B33–B44.

Brown, R. (1973). *A first language: The early stages*. Cambridge, MA: Harvard University Press.

Conway, C. M., & Christiansen, M. H. (2005). Modality-constrained statistical learning of tactile, visual, and auditory sequences. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 31*, 24–39.

Dell, G. S., Reed, K. D., Adams, D. R., & Meyer, A. S. (2000). Speech errors, phonotactic constraints, and implicit learning: A study of the role of experience in language production. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 26*, 1355–1367.

Dienes, Z., Altmann, G. T. M., Kwan, L., & Goode, A. (1995). Unconscious knowledge of artificial grammars is applied strategically. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 21*, 1322–1338.

Eimas, P. D., & Miller, J. L. (1992). Organization in the perception of speech by young infants. *Psychological Science, 3*, 340–345.

Eimas, P. D., Siqueland, E. R., Jusczyk, P., & Vigorito, J. (1971). Speech perception in infants. *Science, 171*, 303–306.

Gebhart, A. L., Aslin, R. N., & Newport, E. L. (submitted for publication). Changing structures in midstream: Learning along the statistical garden path.

Gleason, J. B., Perlmann, R. Y., & Greif, E. B. (1984). What's the magic word? Learning language through routines. *Discourse Processes, 6*, 493–502.

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

Gómez, R. L., Bootzin, R. R., & Nadel, L. (2006). Naps promote abstraction in language learning infants. *Psychological Science, 17*, 670–674.

Hauser, M. D., Newport, E. L., & Aslin, R. N. (2001). Segmentation of the speech stream in a nonhuman primate: Statistical learning in cotton-top tamarins. *Cognition, 78*, B53–B64.

Hockema, S. A. (2006). Finding words in speech: An investigation of American English. *Language Learning and Development, 2*(2), 119–146.

Jenkins, J. G., & Dallenbach, K. M. (1924). Obliviscence during sleep and waking. *American Journal of Psychology, 35*, 605–612.

Johnson, E. K., & Jusczyk, P. W. (2001). Word segmentation by 8-month-olds: When speech cues count more than statistics. *Journal of Memory and Language, 44*, 548–567.

Jusczyk, P. W., Houston, D. M., & Newsome, M. (1999). The beginnings of word segmentation in English-learning infants. *Cognitive Psychology, 39*, 159–207.

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

Kelly, M. H., & Bock, J. K. (1988). Stress in time. *Journal of Experimental Psychology: Human Perception and Performance, 14*, 389–403.

Kinder, A., & Assmann, A. (2000). Learning artificial grammars: No evidence for the acquisition of rules. *Memory and Cognition, 28*, 1321–1332.

Kirkham, N. Z., Slemmer, J. A., & Johnson, S. P. (2002). Visual statistical learning in infancy: Evidence for a domain general learning mechanism. *Cognition, 83*, B35–B42.

Kuhl, P. K. (2004). Early language acquisition: Cracking the speech code. *Nature Reviews Neuroscience, 5*, 831–843.

Kuhl, P. K., Williams, K. A., Lacerda, F., Stevens, K. N., & Lindblom, B. (1992). Linguistic experience alters phonetic perception in infants by 6 months of age. *Science, 255*, 606–608.

Maquet, P., Peigneux, P., Laureys, S., & Smith, C. (2002). Be caught napping: You're doing more than resting your eyes. *Nature Neuroscience, 5*, 618–619.

Marian, V., & Spivey, M. (2003). Competing activation in bilingual language processing: Within- and between-language competition. *Bilingualism: Language and Cognition, 6*, 97–115.

Mattys, S. L., Jusczyk, P. W., Luce, P. A., & Morgan, J. L. (1999). Phonotactic and prosodic effects on word segmentation in infants. *Cognitive Psychology, 38*, 465–494.

Meulemans, T., & Van der Linden, M. (1997). Associative chunk strength in artificial grammar learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 23*, 1007–1028.

Meuter, R., & Allport, D. (1999). Bilingual language switching in naming: Asymmetrical costs of language selection. *Journal of Memory and Language, 40*, 25–40.

Newport, E. L., & Aslin, R. N. (2004). Learning at a distance 1. Statistical learning of non-adjacent dependencies. *Cognitive Psychology, 48*, 127–162.

Odlin, T. (1989). *Language transfer: Cross-linguistic influence in language learning*. Cambridge: Cambridge University Press.

Onishi, K. H., Chambers, K. E., & Fisher, C. (2002). Learning phonotactic constraints from brief auditory experience. *Cognition, 83*(1), B13–B23.

Perruchet, P., & Pacton, S. (2006). Implicit learning and statistical learning: One phenomenon, two approaches. *Trends in Cognitive Sciences, 10*, 233–238.

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

Plihal, W., & Born, J. (1997). Effects of early and late nocturnal sleep on priming and spatial memory. *Psychophysiology, 36*, 571–582.

Rodriguez-Fornells, A., Rotte, M., Heinze, H., Noesselt, T., & Münte, T. (2002). Brain potential and functional MRI evidence for how to handle two languages with one brain. *Nature, 415*, 1026–1029.

Rodriguez-Fornells, A., van der Lugt, A., Rotte, M., Britti, B., Heinze, H. J., & Münte, T. F. (2005). Second language interferes with word production in fluent bilinguals: Brain potential and functional imaging evidence. *Journal of Cognitive Neuroscience, 17*, 422–433.

Saffran, J. R. (2002). Constraints on statistical language learning. *Journal of Memory and Language, 47*, 172–196.

Saffran, J. R., Aslin, R. N., & Newport, E. L. (1996). Statistical learning by 8-month-old infants. *Science, 274*, 1926–1928.

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.

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.

Seidenberg, M. S., MacDonald, M. C., & Saffran, J. R. (2002). Does grammar start where statistics stop? *Science, 298*, 553–554.

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

Spivey, M., & Marian, V. (1999). Cross talk between native and second languages: Partial activation of an irrelevant lexicon. *Psychological Science, 10*, 281–284.

Stickgold, R., & Walker, M. (2005). Memory consolidation and reconsolidation: What is the role of sleep. *Trends in Neurosciences, 28*, 408–415.

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.

Thiessen, E. D., & Saffran, J. R. (2007). Learning to learn: Infants' acquisition of stress-based strategies for word segmentation. *Language Learning and Development, 3*, 73–100.

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

Toro, J. M., Sinnett, J. B. S., & Soto-Faraco, S. (2005). Speech segmentation by statistical learning depends on attention. *Cognition, 97*, B25–B34.

Toro, J. M., & Trobalón, S. J. B. (2005). Statistical computations over a speech stream in a rodent. *Perception and Psychophysics, 67*, 867–875.

Turk-Browne, N. B., Jungé, J., & Scholl, B. J. (2005). The automaticity of visual statistical learning. *Journal of Experimental Psychology, 134*, 552–564.

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

Weber, A., & Cutler, A. (2006). First-language phonotactics in second-language listening. *Journal of the Acoustical Society of America, 119*, 597–607.

Werker, J. F., & Lalonde, C. E. (1988). Cross-language speech perception: Initial capabilities and developmental change. *Developmental Psychology, 24*, 672–683.

Werker, J. F., & Tees, R. C. (1984a). Cross-language speech perception: Evidence for perceptual reorganization during the first year of life. *Infant Behavior and Development, 7*, 49–63.

Werker, J. F., & Tees, R. C. (1984b). Phonemic and phonetic factors in adult cross-language speech perception. *Journal of the Acoustical Society of America, 75*, 1866–1878.

Werker, J. F., & Tees, R. C. (2005). Speech perception as a window for understanding plasticity and commitment in language systems of the brain. *Developmental Psychobiology, 46*(3), 233–251.

Werker, J. F., & Yeung, H. (2005). Infant speech perception bootstraps word learning. *Trends in Cognitive Sciences, 9*, 519–527.

Yang, C. D. (2004). Universal grammar, statistics or both? *Trends in Cognitive Sciences, 8*(10), 451–456.