The influence of phonotactics on suffix discovery in infancy

Liang, Kevin

2024

UNIVERSITY OF CALIFORNIA

Los Angeles

The influence of phonotactics on suffix discovery in infancy

A thesis submitted in partial satisfaction
of the requirements for the degree Master of Arts
in Linguistics

by

Kevin Liang

2024

ABSTRACT OF THE THESIS

The influence of phonotactics on suffix discovery in infancy

by

Kevin Liang

Master of Arts in Linguistics
University of California, Los Angeles, 2024
Professor Megha Sundara, Chair

Recent research has shown that 6-month-olds relate novel words suffixed with -s like *babs* and *teeps* that are embedded in passages, with just the stem *bab* and *teep*, demonstrating an early sensitivity to morphological relatedness. This study explores the limits of early morphological decomposition and its interaction with phonotactics. In two experiments, we evaluated whether monolingual English-learning 6-month-olds are sensitive to phonotactic well-formedness when detecting morphological relatedness. In Experiment 1, we tested infants on two different allomorphs of the English *-s* suffix: [-s] and [-z]. Then, in Experiment 2, we investigated whether infants decompose two different kinds of CV[z] sequences - one type where decomposition leads to stems with permissible sequences in English and another that creates stems with sequences that are not permissible.

Our results indicate several important findings. Firstly, infants possess detailed allomorph-specific representations early in language development with the frequency of different allomorphs influencing early morpheme decomposition. Infants show sensitivity to the more common [-z] allomorph but not the less common [-s] allomorph. Secondly, early morpheme decomposition is not mandatory; infants do not segment every sequence containing a morpheme. Finally, phonotactic cues, even in the absence of semantic or distributional ones, play a crucial role in infants' morphological learning processes. These experiments highlight the interaction between phonotactic sensitivity and morphological acquisition from the earliest stages of language development.

The thesis of Kevin Liang is approved.

Laurel L. Perkins

Bruce P. Hayes

Megha Sundara, Committee Chair

University of California, Los Angeles

2024

献给我的家人

Dedicated to my family

Contents

List of Figures

List of Tables

1. Introduction

An important task that any child acquiring their native language needs to accomplish is the discovery of their language's morphemes - the smallest meaningful units in a language. For example, to acquire English, a child needs to learn that when the morpheme *-s* is attached to a verb, it marks number (singular), person (third), and tense (present). Learning this morpheme is crucial to the acquisition process as it represents a basic building block for developing later syntax.

The process of learning morphology begins very early in infancy. Even before infants utter their first words, they already begin to exhibit some awareness of the morphemes of their native languages. We know that 7.5-month-old English-learning infants can use the morpheme -*ing* to segment frequently occurring, familiar words from running speech (Willits et al., 2014). More recent work has shown that this morphological awareness begins as early as 6-months and extends to entirely novel words. Specifically, Kim & Sundara (2021) show that 6-month-old English-learning infants are able to relate a novel word suffixed with -s like *babs* and *dops* with the corresponding stem *bab* and *dop.*

How do preverbal infants begin to develop this sensitivity to morphology? Looking at other domains of language acquisition, there are several different sources of information that the young learner may exploit to discover morphemes. These include information in the sound segments and sequences themselves, referred to as phonological cues (e.g., Gerken, Wilson, & Lewis, 2005 in learning syntactic categories; e.g., Mattys et al., 1999 in word segmentation), the meanings of such sequences (Braine et al., 1990 in learning word classes) as well as their surrounding contexts, referred to as distributional cues (e.g., Reeder, Newport, & Aslin, 2009). In the experiments discussed in this paper, we begin to disentangle how these cues may be used by infants to decompose sequences into morphemes.

Previous work has only investigated the role of various cues in other areas of language acquisition. No research to date has investigated the role of specific cues in the discovery of morphemes by infants. However, a few studies have shown the ability of adults to employ both distributional and phonological cues, specifically those involving restrictions on permissible sequences of sounds in a language (also referred to as phonotactics), when decomposing a novel word into morphemes. Using an artificial grammar learning (AGL) paradigm, Finley & Newport (2010) demonstrate that adults can use distributional information from frequently recurring final syllables to decompose novel trisyllabic words into stem+suffix sequences. Specifically, they trained adults using audio recordings of 24 CVCV sequences followed by four different CV suffixes. After training, the participants were asked to decide which of two novel suffixed sequences were in the language they were trained on. For example, a participant familiarized with the forms *nobo-ke* and *tise-bu* was presented with the forms *nobobu* (novel but follows the suffixation pattern) and *nokebo* (novel and does not follow the suffixation pattern) during testing. The participants correctly identified the forms which followed the suffixation pattern as being in the novel language (as with *nobobu*) above chance, showing that they were able to decompose the original novel CVCVCV sequences into stem+suffix pairs. Given the use of an AGL paradigm, the participants did not have access to semantic cues nor language-specific phonological cues and could only rely on the distributions of the nonce suffixes.

There is also evidence that adults can take advantage of phonotactic cues to find boundaries of words and morphemes. In a series of word-spotting tasks, Norris et al. (1997) found that adult English speakers had significantly more trouble in identifying the word *apple* in the context *fapple* compared to the context *vuffapple*. They claim that this is because segmenting *fapple* into *f+apple* results in stranding a single consonant *f* which cannot be a possible word of English as it violates the Possible Word Constraint in English (Norris et al., 1997) which states that words in English cannot be a single consonant. On the other hand, segmenting *apple* from *vuffaple* leaves *vuff* which can be a possible word of English as *vuff* does not violate the Possible Word Constraint.

Other phonotactic cues such as restrictions on the placement of vowels can also be used by adults when placing word and morpheme boundaries. In English, there are no words that end in lax vowels (a class of sounds in American English which includes [ɪ] as in *kit*, [ɛ] as in *dress*, [æ] as in *trap*, [ʊ] as in *foot,* and [ʌ] as in *strut*; Hayes, 2011). Using a phrase-picture matching task, Skoruppa et al. (2015) show that English-speaking adults can use the lax vowel constraint to insert word boundaries. In this experiment, participants were told that they were going to hear novel adjective-noun sequences to describe a set of alien creatures. First, they were presented with a novel three-syllable sequence within a carrier phrase (e.g., *This is a* [naɪvʊʃaʊ]) along with a visual representation of a multi-coloured alien creature. Then, the participants heard a new sequence (e.g. *Where is the* [naɪzʌteɪ]?) and were asked to select between three drawings of different aliens - one drawing of a different alien with the same colour (corresponding to a decomposition of [naɪzʌteɪ] into [naɪ] + [zʌteɪ] as [naɪ] occurs in the original sequence [naɪvʊʃaʊ] but [zʌteɪ] does not), one drawing of a different alien in a different colour (corresponding to a decomposition of [naɪzʌ]+[teɪ] as neither [naɪzʌ] nor [teɪ] occur in [naɪvʊʃaʊ]), and a control item. They found that the participants were far more likely to parse the sequence such that there were no lax vowels at the end of the sequence, dispreferring parses like [naɪzʌ]+[teɪ]. Because the participants were presented with nonce words, they did not have access to cues from known lexical items, showing that phonotactic cues are a particularly useful source of information in morpheme decomposition in the absence of meaning.

Converging evidence that adults can learn phonotactics easily and fluently with minimal input comes from studies on non-Māori-speaking New Zealanders. Despite most non-Māori-speaking New Zealanders only explicitly knowing a very limited number of Māori words, they give well-formedness ratings of nonce words that are almost indistinguishable from those of native Māori speakers (Oh et al., 2020; Panther et al., 2023; Panther et al., 2024). These situations wherein speakers develop phonotactic knowledge with limited access to meaning are precisely the situations that infants face in the early stages of acquisition.

Restrictions on the possible sequences of segments within a single morpheme are also used by adults in morpheme decomposition. Hay (2001) presented native English speakers with pairs of novel words such as *vilfim* and *vipfim*. The participants were then asked to select which word was more "complex" (i.e., decomposable). These pairs of words differed in that half of the words contained sequences that were phonotactically illegal in a single morpheme. Consider the novel words *vilfim* and *vipfim* used in the experiment. The sequence [pf] as in *vipfim* never occurs within a single morpheme in English; thus, it is phonotactically ill-formed unless a listener places a morpheme boundary between *vip* and *-fim*. On the other hand, [lf] as in *vilfim* occurs morpheme-internally in many English words (e.g., *dolphin, alpha, sulphur,* etc.). This means that, based on the phonotactics alone, speakers cannot definitively ascertain whether or not to place a morpheme boundary in words with such sequences. Hay (2001) found that it was precisely those novel words that contained illicit morpheme-internal sequences such as [pf] that participants rated as being more complex. As the participants were only asked for judgements on entirely novel stems and suffixes, they did not have access to meaning. These results demonstrate that adult speakers are also sensitive to phonotactic restrictions about sequences within a particular stem or morpheme.

Overall, we see that there are several types of phonotactic cues that are available to adults that can facilitate morpheme decomposition. Firstly, there are word minimality constraints such as the Possible Word Constraint. The role of this constraint in morpheme decomposition can be seen by the fact that adult native speakers have greater difficulty in performing a task that requires positing a boundary that leads to a sub-minimal word (Norris et al., 1997). There are also segmental phonotactic constraints that restrict the presence of particular segments in certain positions. For example, in English, lax vowels cannot occur word-finally, and adult speakers are less likely to insert boundaries that will create such illicit sequences (Skoruppa et al., 2015). Lastly, we also see that there are junctural phonotactic constraints that place restrictions on the sequences of segments that can occur within a single morpheme. For example, in English, [fh] (as in *self-help*) cannot occur within a single morpheme. Hay & Baayen (2003) show that such phonotactically illicit junctures are correlated with both the productivity and parsability (or decomposability) of an affix.

In this paper, we investigated the potential contribution of these types of phonotactic cues in early morphological decomposition. To do so, we conducted two experiments with

English-learning 6-month-olds to explore infants' emerging sensitivity to phonotactics and its interaction with the acquisition of morphology. By 6-months, infants are starting to become sensitive to phonotactics (for a meta-analysis, see Sundara et al., 2022), making it an ideal age to investigate the interaction of phonotactics with morphological decomposition.

In Experiment 1, we investigated the ability of infants to relate CVC[z] (e.g., *babs*) and CVC[s] (e.g., *teeps*) sequences with their CVC stems. These two sequences correspond to two possible realizations of the English -*s* suffix. When the -*s* suffix occurs after a voiced sound (those produced with vibrations of the vocal cords), it is pronounced as [z] such as in *runs* or *hides*, whereas when -*s* occurs after a voiceless sound (those produced without vibration of the vocal cords), it is pronounced as [s] such as in *eats* or *hits*. Importantly, these two sequences also differ in their ability to occur as suffixed and unsuffixed forms. In English infant-directed speech, CVC[z] sequences exclusively occur as suffixed forms such as *runs* and *hides* (see section 2). That is, in infant-directed speech, CVC[z] sequences have absolute cues to morpheme decomposition because they are phonotactically ill-formed without a morpheme boundary. In contrast, CVC[s] sequences can occur as both suffixed (e.g., *kick-s*) and unsuffixed sequences (e.g., *box*). Thus, CVC[s] sequences at best only contain probabilistic, that is, gradient cues to decomposition. Then, in Experiment 2, we tested infants on two different kinds of CV[z] sequences which varied in the phonotactic well-formedness of their CV stems. In one condition, the sequences contained tense vowels (e.g., *geez* [giz] rhymes with *bees*), whereas, in the other condition, they contained lax vowels (e.g., *gihz* [gɪz] rhymes with *fizz*), leading to phonotactically ill-formed stems (e.g., *gih* [gɪ] by itself violates the English lax vowel constraint). Because we used nonce targets in all experiments, infants had no access to semantic cues. We further controlled the distributional information from the surrounding sentence context, so that they were comparable in all cases.

Through these experiments, we sought to address several questions. Are early morphological representations restricted to particular allomorphs? How do gradient and absolute phonotactic cues affect the decomposition of different allomorphs? Is morphological decomposition obligatory in the early stages of acquisition? In other words, once the learner discovers a morpheme, do they decompose every sequence that possibly contains that morpheme? Or are infant learners sensitive to stem phonotactics when performing morpheme decomposition? We use the results of these experiments to examine the contribution of phonotactics in the learning of morphology.

2. Experiment 1: Do infants represent allomorphs?

Kim & Sundara (2021) demonstrated that English-learning 6-month-olds could relate novel words suffixed with *-s* with just the bare stem, but their experimental design did not distinguish the two different allomorphs present in the stimuli: [-s] and [-z]. Four CVC nonce target words were used in their experiment: *bab, kell, dop,* and *teep* (/bæb/, /dɑp/, /kɛl/, and /tip/, respectively). Of these target words, two of the final consonants were voiced (*bab, kell*) while two were voiceless (*dop, teep*). These different final consonants meant that there were two different possible realizations of the English suffix *-s* for the nonce words - [z] when the final consonant was voiced as in *bab* and *kell*, and [s] when the final consonant was voiceless as in *dop* and *teep*. In their experiments, half of the infants were familiarized with *babs* and *dops* while the other half were familiarized with *kells* and *teeps*, resulting in all the infants being exposed to a mixture of both the [-s] and [-z] allomorphs. Thus, from their results, it is unclear whether 6-month-olds could relate forms suffixed with the [-z] as well as the [-s] allomorph.

Experiment 1 was designed to determine whether 6-month-old English-learning infants have discovered both allomorphs of *-s* or just one. In Experiment 1, we familiarized infants with two different sets of sequences - *babs* and *kells* or *dops* and *teeps*. Thus, half of the infants heard sequences containing the [-z] allomorph (*babs, kells*), whereas the other half heard sequences containing the [-s] allomorph (*dops, teeps*). Then, infants were presented with all four stems (*bab, kell*, *dop, teep*) in the test phase. If the infants do not decompose the potentially suffixed sequences (e.g., *babs, kells*), then all four stems (*bab, kell*, *dop, teep*) should be novel and there should be no differences in their listening times between *bab, kell* and *dop, teep*. On the other hand, if the infants do decompose the potentially suffixed sequences (e.g., *babs, kells*), then two of the stems should be familiar (e.g., *bab, kell*) and two should be novel (e.g., *dop, teep*). Thus, if infants listen significantly longer to stems of the suffixed nonce words presented in the familiarization phase compared to completely novel stems, we can infer successful morpheme decomposition. With respect to the [-s] and [-z] conditions, if early morphological representations are restricted to individual allomorphs, then we expect the infants to perform differently in the two conditions. However, if infants begin to develop sensitivity to both the [-s] and [-z] allomorphs at a similar time, then we expect that they will succeed in both conditions.

The results from Experiment 1 also served as a baseline for Experiment 2 by establishing whether English-learning 6-month-olds could segment nonce words suffixed with the allomorph [-z] as [-z] was the only allomorph of *-s* used in the stimuli of Experiment 2.

A priori, there are four reasons that we may expect infants to successfully relate both CVC[s] and CVC[z] sequences and their stems.

Firstly, it may be possible that infants are using overlap in meaning or word class to discover morphemes (Baayen et al., 2016), which could allow them to successfully relate suffixed novel words and their stems in both the [-s] and [-z] conditions. However, this approach is unlikely at 6-months. 6-month-olds have very limited vocabularies (Bergelson & Swingley, 2012). Moreover, the novel items in our experiment were used as verbs, and early vocabularies of infants rarely include verbs, for which it is much harder for infants to assign meaning (Golinkoff & Hirsh-Pasek, 2006). Nor is there any evidence that infants have access to information about word class till the second year of life (Gerken et al., 2005). Additionally, in the present study, semantic cues were unavailable to the infants as all the items were novel words.

Secondly, infants may be able to use distributional information from the sentence context surrounding a target word to determine whether a word is suffixed or unsuffixed. Under such an approach, the learner may keep track of *frequent frames* (Mintz, 2003) of the form A X B where X represents a word and A and B represent the words immediately preceding and following X, respectively. Among these frames, there may be particularly informative ones (e.g., C Y D) where the vast majority of intervening forms are suffixed. For example, most words occurring in the frame "mama/mommy _ a" are suffixed in infant-directed speech. Thus, once a learner hears the context C _ D, they can guess with fairly high accuracy that the intervening form contains a morpheme boundary. Such contexts cannot distinguish between words suffixed with [-s] and words suffixed with [-z], instead, they only provide cues as to whether an intervening form is suffixed or unsuffixed. This means that, if infants are sensitive to this type of distributional information, then we expect that they will succeed in both [-s] and [-z] conditions. Once again, although this approach may be a plausible mechanism for morpheme decomposition, our passages were such that frames were uninformative as to whether the intervening word was suffixed or unsuffixed (see 2.1 for details). As such, the infants could not rely on distributional cues in the present study.

A third possibility is that infants relate any two sequences with sufficient overlap in form. In other words, we might predict that a child will relate *babs* and *bab* due to the fact that *bab* is present at the beginning of both sequences. If overlap in the onset of stems and suffixed forms is enough for infants to relate two forms, then we expect that infants will successfully relate both CVC[s] and CVC[z] sequences with their stems. This is because there is a comparable amount of overlap in both conditions as all the suffixed forms share their initial CVC sequences with their stems. Although plausible, this seems unlikely. Kim & Sundara (2021) show that 6-month-old English-learning infants fail to relate *babbing or babbed* or *babsh* and *bab* even though they successfully relate *babs* and *bab*. Thus, it is likely that phonological overlap alone is not enough for a child to relate two forms.

Lastly, it is not obvious whether 6-month-olds are even able to discriminate between [s] and [z] in word-final position. Fais et al. (2009) demonstrate that 6-month-old English-learning children can distinguish CVC and CVC[s] sequences. Therefore, 6-month-old English-learning infants can - at a minimum - differentiate forms suffixed with *-s* from their stems. However, the discrimination of voicing contrasts in codas such as with CVC[s] and CVC[z] sequences has been shown to be particularly challenging for young children. At 6-months, English-learning infants can only distinguish V[z] from V[s] sequences when both vowel length and voicing cues are present (Eilers et al., 1977). Even adult English speakers exhibit difficulties in distinguishing the voicing contrasts in codas (Chong & Garellek, 2018) as English voiced consonants are often devoiced or weakly voiced. If infants cannot distinguish CVC[s] from CVC[z] sequences, then we expect them to relate CVC stems to both CVC[s] and CVC[z] suffixed forms.

Conversely, there are also many reasons to expect that infants - at least when they begin to discover morphemes - will only successfully relate CVC[z] sequences with their stems but not CVC[s] sequences. Firstly, the [-z] allomorph is more frequent, that is it combines with more unique words than the [-s] allomorph in infant-directed speech as shown in Table 1. The data from Table 1 are drawn from the Brent corpus (Brent & Siskind, 2001) - a 0.5 million word corpus of speech from 8 American English-speaking mothers directed to infants.

Allomorph	**Word Class**	**Function**	**Type frequency**
[-z]	Verb	3rd singular present	148
	Noun	possessive	183
	Noun	plural	547
Total [-z]			*878*
[-s]	Verb	3rd singular present	100
	Noun	possessive	25
	Noun	plural	217
Total [-s]			*342*
[-ɪz]	Verb	3rd singular present	27
	Noun	plural	66
Total [-ɪz]			*93*

Table 1. Frequency of the allomorphs of *-s* in the ~0.5 million word Brent corpus

Kim & Sundara (2021) argue that the morphemes with the highest type frequency regardless of meaning, function, or word class are discovered the earliest by infants. This is evidenced by the fact that at 6-months, English-learning infants only relate forms suffixed with *-s* - the most frequent suffix in infant-directed speech - and their stems. However, by 8-months, they also begin to relate forms suffixed with *-ing* - the second most frequently occurring suffix in infant-directed speech - with stems. These results support an acquisition trajectory where the discovery of morphemes is consistent with type frequency early on. As the [-z] allomorph occurs with more than double the number of words (*n* = 878) compared to the [-s] allomorph (*n* = 342) in the Brent corpus, if infants begin to develop sensitivity to the morphemes with the highest type frequency regardless of meaning or function, then infants should successfully relate forms suffixed with [-z] with their stems.

Additionally, across all occurrences of [s] and [z], [z] is far more likely to be a morpheme than [s] is. In the Brent corpus, 70% of word-final [z]'s are morphemes. Thus, simply by inserting a morpheme boundary before every occurrence of [z] word-finally, the learner can achieve a high accuracy in morpheme decomposition. However, the reverse is true of only a minority of word-final [s]'s. Word-final [s]'s in the Brent corpus only occur as a morpheme 14% of the time.

Besides the higher type frequency of the [-z] allomorph over the [-s] allomorph and the higher likelihood of [z] being a morpheme, there is another potential reason why we might expect infants to relate CVC[z], but not CVC[s] sequences with their stems, which involves the difference in phonotactics mentioned previously. In English infant-directed speech, sequences such as [bz] and [lz] as in *ba**bs*** and *ke**lls*** are phonotactically ill-formed unless there is a morpheme boundary between the two segments. This means that if the learner does not decompose [bæbz] and [kɛlz] into stem+suffix sequences ([bæb+z] and [kɛl+z]), they are left with a phonotactically illegal sequence. On the other hand, CVC[s] sequences are well-formed both as unsuffixed (e.g. *box*) and suffixed sequences (*e.g. kick+s*). Therefore, an infant may treat CVC[s] sequences as either suffixed or unsuffixed. That is, CVC[z] sequences in infant-directed speech have absolute cues to morpheme decomposition in that they obligatorily signal morpheme decomposition, whereas CVC[s] sequences only offer gradient cues to decomposition as both suffixed and unsuffixed forms are possible.

This distinction can also be seen in an analysis of the Brent corpus. Across the corpus, there are 85 different CVC[z] sequences that occur. All of these sequences are suffixed forms that can be decomposed into CVC+[z] sequences. Since all 85 CVC[z] sequences are suffixed, a learner who decomposes every CVC[z] sequence into two morphemes will achieve a high accuracy. However, the same is not true with CVC[s] sequences. In the Brent corpus, there are 54 suffixed CVC[s] sequences and 17 unsuffixed ones, making the decomposition of such sequences more uncertain.

Indeed, there is evidence that infants behave differently when presented with absolute and gradient phonotactic cues, at least for word segmentation, early on in acquisition. Mattys & Jusczyk (2001) show that English-learning 9-month-olds successfully extract nonce words from running speech when embedded in contexts with absolute phonotactic cues where failing to segment the word would create illegal or low frequency CC clusters. However, they fail when the nonce words are embedded in contexts with gradient phonotactic cues where a word boundary is optional. For example, in the utterance *gaffe hold*, failure to insert a word boundary between *gaffe* and *hold* would lead to a word-internal [fh] cluster which is improbable in English. On the other hand, in the utterance *gaffe tine*, failure to insert a word boundary between *gaffe* and *tine* would lead to a [ft] cluster which occurs both between and within words in English.

In summary, if infants primarily rely on semantic cues, distributional cues, and phonological overlap to perform morpheme decomposition, then we expect their behaviour with respect to the two allomorphs of *-s* to be identical. If, instead, they are sensitive to the higher frequency of the [-z] allomorph over the [-s] allomorph, the higher likelihood of [z] being a morpheme, and phonotactic cues, then we expect that they will successfully decompose CVC[z] sequences but not CVC[s] sequences.

2.1 Stimuli

The stimuli used in Experiment 1 were identical to those used in Kim & Sundara (2021) except the two sub-conditions were re-arranged such that the participants in one sub-condition would only be familiarized on nonce forms suffixed with the [-z] allomorph (*babs* and *kells*) and those in the other sub-condition would only be familiarized on forms suffixed with the [-s] allomorph (*dops* and *teeps*). In total four different suffixed nonce words were used in the experiment: *babs, kells, dops,* and *teeps*. These words were constructed to have different vowel qualities. Moreover, their phonotactic probabilities were calculated using the Phonotactic Probability Calculator (Vitevitch & Luce, 2004) to ensure that the phonotactic probabilities of the nonce words were comparable to existing English words.

The suffixed nonce words were embedded in four six-sentence passages where each occurrence of the suffixed form was preceded by *mommy* or *mama* - a common word known to help with word segmentation (Bortfeld et al., 2005). The use of *mommy* and *mama* in the passages was to help the infants attune to the suffixed nonce forms in the familiarization phase. We also recorded four lists containing just the stems in isolation - *bab, kell, dop,* and *teep* - for use in the testing phase.

To investigate the viability of distributional cues to morpheme decomposition in the experiment, we examined the surrounding sentence contexts used in the passages for the experiment. Table 2 shows all of the frames used in the passages (Appendix A) of the form A _ B where _ represents the position of the suffixed nonce word and A and B represent the words immediately preceding and following the nonce word, respectively. We counted the number of types (distinct stems) of suffixed and unsuffixed words that occur in the context A _ B in the Brent corpus (Brent & Siskind, 2001). Those frames that occur in the passages but not in the Brent corpus are excluded from the table.

Frame	Unsuffixed	Suffixed	Suffixed ratio
Mama/Mommy _ a	11	14	56%
Mama/Mommy _ and	7	4	36%
Mama/Mommy _ if	1	0	0%
Mama/Mommy _ # (end of utterance)	99	33	25%
Mama/Mommy _ (bigram, one-sided frame)	225	146	39%

Table 2. Frequency of suffixed and unsuffixed sequences between frames in the Brent corpus

The majority of items in infant-directed speech which occur in the same frames as the ones used for the nonce words in the experiment, are unsuffixed forms. Even for the best frame

(Mama/Mommy _ a), the ratio of suffixed sequences is still only 56%. Thus, due to the particular passages used in the experiment, distributional information from the surrounding sentence context is unlikely to be useful to the learner in determining whether an intervening form is suffixed or not.

The stimuli were recorded by a 25-year-old female native English speaker from Southern California who was unfamiliar with the purpose of the experiment. She was instructed to read the passages in an animated voice as if she were talking to a preverbal infant. In total, six four-sentence passages containing the suffixed nonce words and four lists containing 15 repetitions of the stems in isolation were used. The passages used for testing can be found in Appendix A. These stimuli are identical to those used in Kim & Sundara (2021), and detailed acoustic analyses of the stimuli are reported there.

2.2 Participants

In total, the data from 60 (30 per condition; 22 female) 6-month-olds (mean = 187 days, range 168-209) were included in the experiment. The infants' language exposure was estimated using a language questionnaire (Sundara & Scutellaro, 2009), and only those infants with at least 90% of their language input as English were included (mean = 98%, range 90-100). Furthermore, based on parental reports, none of the infants had a history of speech, language, or hearing difficulties, and were in good health on the day of testing. Additional infants were tested but not included in the final dataset because they did not complete testing due to fussiness (*n* = 8), failure to look at the lights (*n* = 4), exceeding the maximum experiment time of 10 minutes (*n* = 5), technical difficulties ($n = 1$), or having listening times more than two standard deviations away from the mean ($n = 1$).

2.3 Procedure & Design

The Headturn Preference Procedure was used to test the infants (Jusczyk & Aslin, 1995; Kemler Nelson et al., 1995). The choice of whether to use the *mommy* or *mama* passages was determined by asking the parent which form was used at home. Testing took about 10 minutes per participant. Each infant sat on their caregiver's lap facing the center panel of a three-sided pegboard booth with lights attached to each of the three sides. Both the parent and experimenter wore headphones with music playing to prevent them from inadvertently influencing the child.

At the beginning of testing, the light on the center panel flashed to draw the infant's attention. Subsequently, a light on one of the side panels began to flash to attract the infant's attention towards that side of the booth. Once the infant was focused on the light, auditory stimuli began to play from a speaker just below the light. The experimenter observed the infant through a video feed and recorded both the direction and duration of the infant's head turns. Presentation of auditory stimuli was completely contingent on infant looking behavior, and thus served as a proxy for listening time - the primary variable of interest in this study.

In the experiment, the infants were first familiarized with two suffixed nonce words (e.g., either *babs, kells* or *dops, teeps*, counterbalanced) embedded in passages until they accumulated 45 seconds of listening time to each of the two suffixed nonce words. Afterward, in the test phase, the infants were presented with familiar and novel, isolated stems (*bab, kell, dop, teep*) in three blocks for a total of 12 trials. As is typical for segmentation studies with natural language stimuli, significantly different listening times to potentially familiar stems compared to novel stems provide evidence that infants successfully related suffixed words to stems.

2.4 Analysis

The listening times were analyzed with linear mixed-effects models through the *lmerTest* package in R. The fixed effects included the between-subjects variable Condition (whether the infants were familiarized with *babs*, *kells* or *dops*, *teeps*), the within-subjects variable Block (1st, 2nd, or 3rd) - to control for the effects of repeated exposures to the isolated nonce words on listening times, and Trial Type (familiar vs. novel) and all their interactions.

The model also included a random intercept for Subject to control for differences in the baseline listening times. As the model failed to converge when slopes for Trial Type by Subject and Block by Subject were included, the final model did not include a random slope. The anova() function was used to evaluate fixed effects against the full model, and planned comparisons were performed using the *emmeans* package in R.

We also fit models on log-transformed listening times, but the qualitative pattern was the same. Thus, only the results of the models fit directly on the raw listening times are reported to allow comparisons to published results.

2.5 Results & Discussion

The listening time data in seconds for Experiment 1 are shown in Figure 1. As expected, there was a significant main effect of Block (*F*-value = 16.7, $p < 0.001$), confirming that listening time reduced across the experiment. There was a significant main effect of Trial Type (*F*-value = 3.99, $p = 0.046$), indicating that there were significantly different listening times between novel and familiar stems. We also expected a significant interaction of Condition and Trial Type (*F*-value = 2.06, $p = 0.15$), but it was not significant. Crucially for our hypotheses, planned comparisons showed that the main effect of Trial Type was significant for the [-z] allomorph condition [*t ratio* = -2.4, $p = 0.02$] but not the [-s] allomorph condition [*t ratio* = 0.27, $p = 0.8$], although the interaction between Condition and Trial Type (*F*-value = 2.06, $p = 0.15$), was not significant.

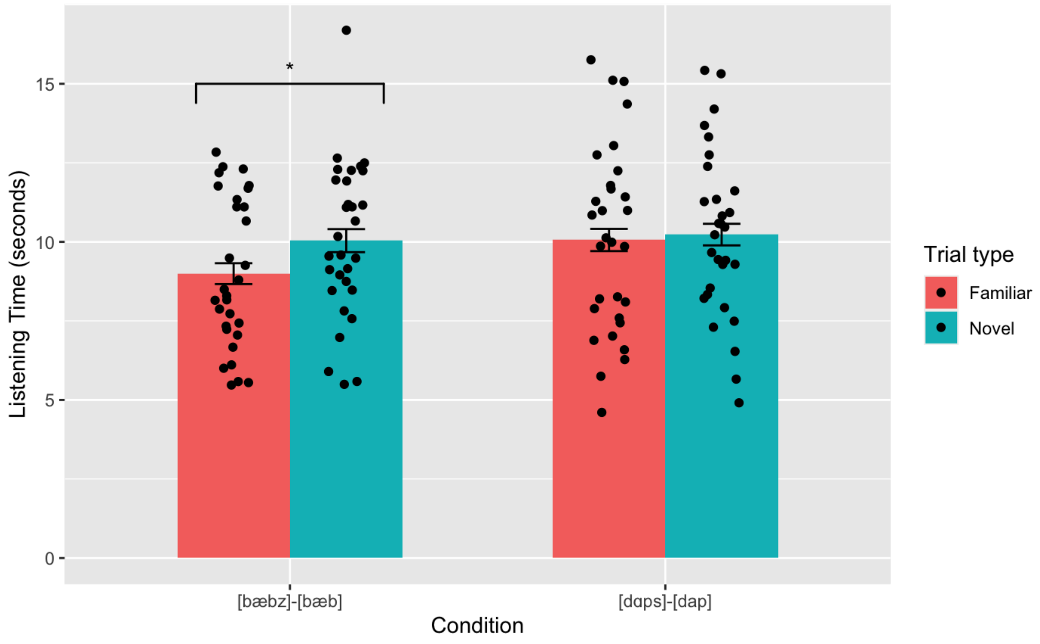

Figure 1. Mean listening times (in seconds) by condition and trial type (Expt. 1)

None of the other main effects or interactions were significant. There was no significant main effect of Condition (F-value = 1.01, p = 0.32). There also was no significant interaction of Block and Condition (F-value = 0.14, p = 0.87), Block and Trial Type (F-value = 0.49, p = 0.61), nor Block, Condition, and Trial Type (F-value = 1.06, p = 0.35).

Thus, infants listened significantly longer to potentially familiar novel stems compared to completely unfamiliar novel stems when familiarized with suffixed forms containing the allomorph [-z] but not [-s]. In other words, 6-month-olds only displayed sensitivity to the [-z] allomorph of the English -*s* suffix as they related suffixed nonce forms such as *babs* with its stem *bab* but not suffixed nonce forms such as *dops* with its stem *dop*.

The results of Experiment 1 are consistent with an account where early morphological representations by infants are allomorph specific. It is precisely [-z] - the allomorph with the higher type frequency - to which 6-month-olds display sensitivity. At the same time, 6-month-olds do not yet display any sensitivity to the [-s] allomorph. The fact that there was evidence only that infants successfully decompose [-z] but not [-s] is incompatible with the idea that infants solely rely on phonological overlap to relate forms; there is comparable overlap between CVC[s] and CVC[z] forms and their stems, yet infants only succeeded in one condition. Recall also that semantic information was unavailable to the infants in the experiment because they were tested on their abilities to segment suffixes from nonce words.

These findings are also compatible with an account where infants as young as 6-months are sensitive to phonotactic cues when performing morpheme decomposition as it is the condition with absolute phonotactic cues (CVC[z]) where the infants succeeded. When the phonotactic cues were gradient (CVC[s]), they failed to decompose morphemes. It is not possible to disentangle the contribution of phonotactics from that of frequency solely based on the results of Experiment 1. To better investigate the role of phonotactics specifically, in Experiment 2, we focus on the [-z] allomorph which infants can successfully decompose.

Finally, because infants were successful at relating CVC[z] with CVC stems but not CVC[s] with the same stems, we can be certain that they can hear the differences between [z] and [s], in final position, in consonant clusters.

3. Experiment 2: Are infants sensitive to stem phonotactics in performing morpheme decomposition?

In Experiment 2, we tested English-learning 6-month-olds on a new set of nonce verbs using the same sentence frames as in Experiment 1. However, unlike in Experiment 1, these nonce verbs had a different shape. Instead of novel CVC[s] or CVC[z] sequences, we familiarized infants on novel CV[z] sequences (such as *gihz* [gɪz] and *geez* [giz]). This experiment sought to address two questions: 1) what are the limits of infant morphological decomposition? Is morpheme decomposition obligatory once an infant discovers a morpheme?; and 2) what is the role of phonotactics in early morpheme decomposition? To address these questions, we familiarized the infants on two different kinds of CV[z] sequences containing either tense vowels or lax vowels.

For the tense vowel sequences, decomposing the CV[z] sequences into a CV stem followed by the suffix [-z] would lead to a phonotactically legal CV stem. Given the success of English-learning 6-month-olds in decomposing novel words suffixed with the [-z] allomorph in Experiment 1, we expected them to successfully decompose $CV_{[+tense]}$[z] sequences (e.g., *geez* to *gee*+*z*).

For the lax vowel sequences, decomposing the CV[z] sequences into stem+suffix sequences would lead to phonotactically illegal stems. These sequences are phonotactically illegal in two different ways. Firstly, the stems produced by decomposing $CV_{[+lax]}$[z] into $CV_{[+lax]}$+[z] sequences (e.g., *gihz* [gɪz] to *gih* [gɪ]) end in lax vowels which violates the English phonotactic restriction against words ending in lax vowels (e.g., *gih* -*[gɪ] and *deh* - *[dɛ] are not possible English words). Secondly, $CV_{[+lax]}$ stems are ill-formed in that they are not Minimal

Prosodic Words (Fee, 1992). This is also referred to as the word minimality constraint and formalizes the observation that languages differ in what constitutes a minimal possible word (for an overview of word minimality cross-linguistically, see Hayes, 1995). In English, all words must have at least two moras - the basic timing unit of which short syllables have one and long syllables have two. All lax vowels in English are short, meaning that they only contain one mora. Thus, $CV_{[+lax]}$ stems only contain a single mora and violate the English word minimality constraint.

If infants obligatorily segment morphemes regardless of phonotactic cues after they first discover a morpheme, then they should successfully relate sequences such as *dehz* [dɛz] with *deh* [dɛ], just like they relate *babs* with *bab*. In other words, an infant learner who solely relies on the frequency of the allomorph would be expected to obligatorily decompose sequences ending in [-z]. Conversely, if young infants are sensitive to phonotactic cues when decomposing morphemes, then they should fail to relate such sequences. As with the previous experiment, significantly different listening times to potentially familiar stems compared to novel stems was used as evidence that infants successfully related suffixed words to stems.

3.1 Stimuli

For both tense and lax vowel stimuli, the suffixed nonce words were embedded in four six-sentence passages. These passages were identical to those used in Experiment 1 except the nonce forms were replaced with the new target words.

The stimuli were recorded by a different phonetically trained native English speaker from Experiment 1 who was unfamiliar with the purpose of the experiment. The speaker for Experiment 2 was a 27-year-old female native speaker of Mainstream American English. She was instructed to read the passages in an animated voice as if she were talking to a preverbal infant. The recordings were done in a soundproof booth using a Shure SM10A head-mounted microphone. All stimuli were digitized at a sampling frequency of 22,050 Hz and 16–bit quantization. The average intensity, duration, and pitch of all of the recordings were equalized using Praat (Boersma & Weenink, 2013) as infants have been found to be very sensitive to low-level acoustic cues (Kim & Sundara, 2021). All the measurements and analyses were done using PRAAT (Boersma & Weenink, 2013), and the average loudness level for stimuli during playback was 75 dB.

3.1.1 Tense Vowel Stems

Four nonce target words were used containing tense vowels: *doyz, fooz, geez, tauz* ([dɔɪz], [fuz], [giz], and [taʊz]) along with their stems *doy, fu, gee,* and *tau* ([dɔɪ], [fu], [gi], and [taʊ]). These words were constructed to have different vowel qualities.

The passages had an average duration of 21.8 s (SD = 0.02) and an average pitch of 240.2 Hz (range 60-596). In processing the passages, the intensity of the passages was normalized to 75 dB. The average duration of each of the nonce words in the passages was 647 ms (SD = 128) and 1030 ms (SD = 114) in the lists. The average pitch of the nonce words in the passages was

220 Hz (SD = 33) and 249 Hz (SD = 47) in the lists. The average intensity of the nonce words in the passages was 73.5 dB (SD = 2.3) and 75.7 dB (SD = 3.1) in the lists.

Passages used for testing the tense vowel sequences can be found in Appendix B.

3.1.2 Lax Vowel Stems

Four additional nonce target words were used containing lax vowels: *gihz, dehz, tuhz,* and *bahz* ([gɪz], [dɛz], [tʊz], and [bæz]) along with their stems *gih, deh, tuh,* and *bah* ([gɪ], [dɛ], [tʊ], and [bæ]). As the stems end in lax vowels, they are phonotactically illegal words in English. These words were also constructed to have different vowel qualities. The speaker who recorded these stimuli had phonetic training and was asked to produce natural lax vowels without lengthening them. We judged that, after considerable practice, our speaker was able to do this successfully.

The passages had an average duration of 20.0 s (SD = 0.02) and an average pitch of 221.5 Hz (range 58-599). In processing the passages, the intensity of the passages was normalized to 75 dB. The average duration of each of the nonce words in the passages was 548 ms (SD = 136) and 801 ms (SD = 185) in the lists. The average pitch of the nonce words in the passages was 226 Hz (SD = 46) and 267 Hz (SD = 55) in the lists. The average intensity of the nonce words in the passages was 73.3 dB (SD = 4.4) and 74.5 dB (SD = 3.0) in the lists.

Passages used for testing the lax vowel sequences can be found in Appendix C.

3.2 Participants

In total, the data from 51 (30 lax vowel stems and 21 for tense vowel stems; 31 female) 6-month-olds (mean = 182 days, range 161-204) were included in the experiment. Testing for the remaining 9 participants is still ongoing. As in Experiment 1, the infants' language exposure was estimated using a language questionnaire (Sundara & Scutellaro, 2009), and only those infants with at least 90% of their language input as English were included (mean = 97, range 90-100). Furthermore, based on parental reports, none of the infants had a history of speech, language, or hearing difficulties, and were in good health on the day of testing. Additional infants were tested but not included in the analyzed data because of fussiness (*n* = 5), failure to look at the lights (*n* = 3), technical difficulties (*n* = 1), or having listening times more than two standard deviations away from the mean (*n* = 1).

3.3 Procedure

The procedure was identical to that of Experiment 1 except for the different stimuli used. In the familiarization phase, the infants were first familiarized with two suffixed nonce words (e.g., either *doyz, fooz* or *geez, tauz* for infants tested on tense vowels and either *gihz, dehz* or *tuhz, bahz* for infants tested on lax vowels) embedded in passages until they accumulated 45 seconds of listening time to each of the two suffixed nonce words. Thus, half of the infants exposed to tense vowel sequences were trained on *doyz* and *fooz* while the other half were trained on *geez* and *tauz*. Likewise, half of the infants exposed to lax vowel sequences were trained on *gihz* and *dehz* while the other half were trained on *tuhz* and *bahz*. Afterward, in the test phase, the infants were presented with familiar and novel, isolated stems (either *doy, fu, gi, tau* or *gih, deh, tuh, bah* depending on whether the child was familiarized on tense or lax vowel sequences) in three blocks for a total of 12 trials. Note that each infant either exclusively heard nonce forms containing tense vowels or lax vowels. No participant heard a mix of the stimuli.

3.4 Analysis

The analyses were identical to those of Experiment 1. The fixed effects included the between-subjects variable Tenseness (whether the infants were familiarized with tense or lax vowel stems), the within-subjects variable Block (1st, 2nd, or 3rd), and Trial Type (familiar vs. novel) and all their interactions.

As with Experiment 1, the model also included a random intercept for Subject. Unlike Experiment 1, the model had an additional random slope for Trial Type by Subject. The anova() function was used to evaluate fixed effects against the full model, and planned comparisons were performed using the *emmeans* package in R.

Additionally, as in Experiment 1, we fit models on log-transformed listening times, but the qualitative pattern was the same. Thus, only the results of the models fit directly on the listening times are reported.

3.5 Results & Discussion

The listening time data for Experiment 2 are shown in Figure 2. As expected, there was a significant main effect of Block (*F*-value = 29.0, $p < 0.001$), confirming that there was a reduction of listening times across trials. There was also a significant interaction of Block and Tenseness (*F*-value = 3.37, $p = 0.04$), showing that there was a difference in the reduction of listening times across trials depending on whether the stimuli contained tense or lax vowels. As with Experiment 1, we also predicted that there would be a significant interaction between Tenseness and Trial Type (*F*-value = 1.30, $p = 0.26$), but this interaction did not turn out to be significant.

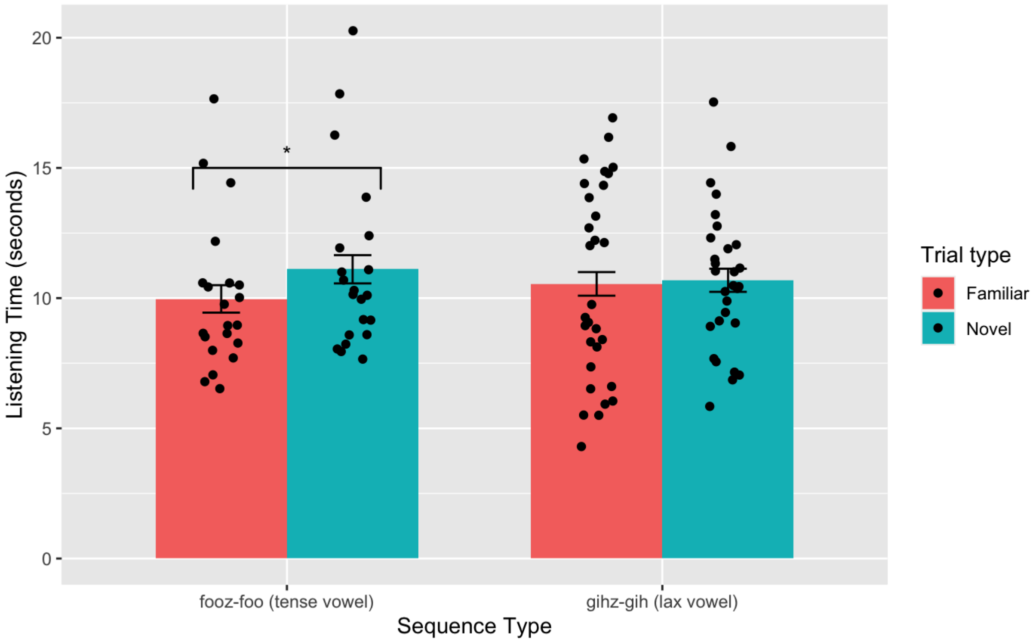

Figure 2. Mean listening times (in seconds) by tenseness and trial type (Expt. 2)

There was also no significant main effect of Tenseness (*F*-value = 0.01, *p* = 0.93) nor Trial Type (*F*-value = 2.10, *p* = 0.15). There also was no significant interaction of Block and Trial Type (*F*-value = 0.51, *p* = 0.60), nor Block, Tenseness, and Trial Type (*F*-value = 2.29, *p* = 0.11).

Planned comparisons showed that the effect of Trial Type was trending towards significance for tense vowel stems [*t ratio* = -1.69, *p* = 0.10] but not lax vowel stems [*t ratio* = -0.25, *p* = 0.81]. Thus, there was no significant difference between the infants' listening times to potentially familiar stems and completely novel stems when the resulting stem ended in a lax vowel and was, thus, phonotactically ill-formed. On the other hand, the infants listened marginally longer to completely novel CV stems compared to potentially familiar CV stems when the resulting stem ended in a tense vowel and was, thus, phonotactically well-formed. In other words, 6-month-olds only related novel suffixed words with their stems if the resulting stem is phonotactically legal.

These results show that, even in the earliest stages of morphological acquisition, infants incorporate phonotactic knowledge when decomposing morphemes. As with Experiment 1, infants could not have relied on distributional or semantic cues because they were not available. This is because the same passages were used for both tense and lax vowel sequences, meaning infants had access to the same distributional information, yet they only succeeded with tense vowels. Moreover, as discussed in 2.1, the distributional cues provided by the passages are not particularly informative of whether or not the nonce form is suffixed. Additionally, as in Experiment 1, the use of nonce items eliminates the possibility of semantic cues to morpheme decomposition.

4. General Discussion

In the present paper, we investigated the ability of monolingual English 6-month-olds to decompose novel words suffixed with *-s* into stem+suffix sequences. In Experiment 1, we showed that the ability of English-learning 6-month-olds to decompose the *-s* suffix was restricted to one particular allomorph - the more frequent allomorph [-z]. They were able to successfully relate CVC[z] sequences with their stems but not CVC[s] sequences.

Having established the capability of infants to decompose the [-z] allomorph in English, we sought to test the limits of this ability. Thus, in Experiment 2, we familiarized infants on two different types of CV[z] sequences and tested whether they could relate them to CV stems. Half of the infants heard sequences containing lax vowels which produced phonotactically ill-formed stems while the other half heard sequences containing tense vowels which produced phonotactically well-formed stems. Ultimately, only the infants tested on tense vowel stems, but not lax vowel stems, succeeded in relating CV[z] suffixed forms with CV stems.

The results of these experiments have several implications for understanding how infants begin to discover morphemes. Firstly, the results of Experiment 1 demonstrate that morphological representations are detailed early on in language development. In fact, English-learning 6-month-olds have distinct representations for the allomorphs [-z] and [-s] as evidenced by their success in decomposing CVC[z] but not CVC[s] sequences. This should perhaps not be surprising if infants have no access to meaning or word class. Additionally, based on this finding we can also infer that at 6-months, English-learning infants can distinguish

CVC[z] from CVC[s] sequences. They do not neutralize the word-final voicing contrast in perception as they behave differently when presented with the two kinds of sequences.

The results of Experiment 1 also demonstrate the importance of frequency in early morpheme decomposition. Infants only related nonce words suffixed with the more frequent allomorph [-z] of the English -s suffix to stems but not nonce words suffixed with the less frequent [-s] allomorph. However, a frequency-based account of morpheme decomposition alone is not adequate. A learner that solely relies on the frequency of the allomorph would be expected to obligatorily decompose sequences ending in [-z]. Such a learner would have succeeded in decomposing CVC[z] sequences and CV[z] sequences regardless of the phonotactic well-formedness of the stem. Instead, infants succeeded in relating CVC[z] sequences and their stems but only successfully related CV[z] sequences and their stems when the stem would be phonotactically well-formed.

This leads us to another important finding which is the fact that, even early in acquisition, morpheme decomposition is not obligatory. The results of both experiments underscore this finding. Even upon the initial discovery of a morpheme, infants do not decompose every sequence that possibly contains that morpheme. Despite successfully decomposing CVC[z] in Experiment 1 and $CV_{[+tense]}$[z] sequences in Experiment 2, English-learning 6-month-olds did not decompose $CV_{[+lax]}$[z] sequences.

The difference in the tendency to decompose the various sequences also shows that phonotactic sensitivity interacts with the learning of morphology even in the earliest stages of acquisition. In Experiment 1, it is precisely the sequence that would be phonotactically illegal if it is treated as unsuffixed (CVC[z]) that is successfully decomposed by 6-month-olds. CVC[s] sequences which can be variably unsuffixed or suffixed are not decomposed. Likewise, in Experiment 2, infants only decomposed CV[z] sequences containing tense vowels but not those containing lax vowels which would lead to phonotactically ill-formed stems.

As the lax vowel stems potentially violate two different phonotactic restrictions - the lax vowel constraint and the English word minimality constraint, the present study is not able to differentiate between the role of each constraint. However, we believe that it is more plausible that 6-month-olds possess some sensitivity to the lax vowel constraint as opposed to word minimality restrictions. This is because acquiring a restriction like the word minimality constraint requires the infant to have representations of subsyllabic structure (such as moras). In a study of the productions of English-learning children, Demuth & Fee (1995) argue that many children do not develop such representations until as late as 2 years. Moreover, at 6-months, infants are still learning to segment words, and we know that even 7.5-month-old English-learning children fail to segment words consisting of an unstressed syllable followed by a stressed syllable (Jusczyk et al., 1999). How infants could have a hypothesis about the minimal length of a word in their native language, without successfully segmenting the entire input, is unclear. In contrast, even if the input is partially segmented, the learner could begin to acquire sensitivity positional restrictions on the placement of lax vowels. They could do so based solely on the absence of lax vowels at word, clause, or utterance boundaries. Regardless of the specific phonotactic cue at play, the fact that the infants failed to decompose a form that would lead to a phonotactically ill-formed stem indicates that they were sensitive to some phonotactic restrictions in their native language.

The role of phonotactics is further supported by the lack of other viable cues in the experiments. As the passages used were comparable across experiments and conditions and the infants were tested on nonce words, the learners did not have access to differing distributional nor semantic information that could explain any differences. Instead, phonotactic cues served as a key difference between the various conditions in this study.

These results are consistent with the adult literature showing that adults use phonotactic cues even in the absence of meaning (Norris et al., 1997; Skoruppa et al., 2015; Oh et al., 2020). Moreover, the effects of phonotactic restrictions within a single stem or morpheme on decomposition can also be observed in adults (Hay, 2001). However, this study is the first to show that infants as young as 6-months are able to incorporate phonotactic knowledge in performing morpheme decomposition.

In summary, monolingual English learning 6-month-olds can relate novel CVC[z], but not CVC[s], sequences with their stems. They also successfully relate CV[z] sequences containing tense vowels with their stems but fail on CV[z] sequences containing lax vowels. Our findings provide evidence for an acquisition trajectory where morphological representations arise early and are sensitive to language-specific phonotactic restrictions.

Appendix A

bab /bæb/

Mommy **babs** and sings at the same time. I feel so happy whenever Mommy **babs**. I play the piano and Mommy **babs**. Mommy **babs** if she sees me dancing around. Mommy **babs** while Grandma and Grandpa eat. If I jump up and down Mommy **babs**.

kell /kɛl/

My daddy always laughs whenever Mommy **kells**. Mommy **kells** a lot and I love it. Grandpa says he smiles because Mommy **kells**. I really like when Mommy **kells**. Mommy **kells** when I play blocks with my brother. Mommy **kells** whenever she is happy.

dop /dɑp/

I get so excited when Mommy **dops**. Mommy **dops** when my brother and I play the drum. Mommy **dops** when she is proud of me. Daddy dances while Mommy **dops**. Mommy **dops** every time she sees me eating. My sister and I jump when Mommy **dops**.

teep /tip/

My brother smiles every time Mommy **teeps**. Mommy teeps whenever she is happy. My sister and I sing and Mommy **teeps**. Mommy **teeps** a lot and so does Daddy. I get so excited when Mommy **teeps**. Mommy **teeps** when I play with my sister.

Appendix B

doys /dɔɪz/

Mommy **doys** and sings at the same time. I feel so happy whenever Mommy **doys**. I play the piano and Mommy **doys**. Mommy **doys** if she sees me dancing around. Mommy **doys** while Grandma and Grandpa eat. If I jump up and down Mommy **doys**.

foos /fuz/

My Daddy always laughs whenever mommy **foos**. Mommy **foos** a lot and I love it. Grandpa says he smiles because Mommy **foos**. I really like when Mommy **foos**. Mommy **foos** when I play blocks with my brother. Mommy **foos** whenever she is happy.

gees /giz/

I get so excited when Mommy **gees**. Mommy **gees** when my brother and I play the drum. Mommy **gees** when she is proud of me. Daddy dances while Mommy **gees**. Mommy **gees** every time she sees me eating. My sister and I jump when Mommy **gees**.

taws /taʊz/

My brother smiles every time Mommy **taws**. Mommy **taws** whenever she is happy. My sister and I sing and Mommy **taws**. Mommy **taws** a lot and so does Daddy. I get so excited when Mommy **taws**. Mommy **taws** when I play with my sister.

Appendix C

gihz /gɪz/

Mommy **gihz** and sings at the same time. I feel so happy whenever Mommy **gihz**. I play the piano and Mommy **gihz**. Mommy **gihz** if she sees me dancing around. Mommy **gihz** while Grandma and Grandpa eat. If I jump up and down Mommy **gihz**.

dehz /dɛz/

I get so excited when Mommy **dehz**. Mommy **dehz** when my brother and I play the drum. Mommy **dehz** when she is proud of me. Daddy dances while Mommy **dehz**. Mommy **dehz** every time she sees me eating. My sister and I jump when Mommy **dehz**.

tuhz /tʊz/

My Daddy always laughs whenever Mommy **tuhz**. Mommy **tuhz** a lot and I love it. Grandpa says he smiles because mommy **tuhz**. I really like when Mommy **tuhz**. Mommy **tuhz** when I play blocks with my brother. Mommy **tuhz** whenever she is happy.

bahz /bæz/

My brother smiles every time mommy **bahz**. Mommy **bahz** whenever she is happy. My sister and I sing and Mommy **bahz**. Mommy **bahz** a lot and so does daddy. I get so excited when Mommy **bahz**. Mommy **bahz** when I play with my sister.

References

Aslin, R., Newport, E., & Reeder, P. (2009). The role of distributional information in linguistic category formation. *Proceedings of the Annual Meeting of the Cognitive Science Society, 31*.

Baayen, R. H., Shaoul, C., Willits, J., & Ramscar, M. (2016). Comprehension without segmentation: a proof of concept with naive discriminative learning. Language, *Cognition and Neuroscience, 31*(1), 106–128. https://doi.org/10.1080/23273798.2015.1065336

Bergelson, E., & Swingley, D. (2012). At 6–9 months, human infants know the meanings of many common nouns. *Proceedings of the National Academy of Sciences, 109*(9), 3253-3258.

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*(4), 298–304.

Braine, M. D., Brody, R. E., Brooks, P. J., Sudhalter, V., Ross, J. A., Catalano, L., & Fisch, S. M. (1990). Exploring language acquisition in children with a miniature artificial language: Effects of item and pattern frequency, arbitrary subclasses, and correction. *Journal of memory and language, 29*(5), 591-610.

Brent, M. R. & Siskind, J. M. (2001). The role of exposure to isolated words in early vocabulary development. *Cognition, 81*(2), 31-44.

Chong, J. A., & Garellek, M. (2018). Online perception of glottalized coda stops in American English. *Laboratory Phonology, 9*(1).

Demuth, K., & Fee, E. J. (1995). *Minimal prosodic words in early phonological development*. Ms., Brown University and Dalhousie University.

Eilers, R. E., Wilson, W. R., & Moore, J. M. (1977). Developmental changes in speech discrimination in infants. *Journal of Speech and Hearing Research, 20*(4), 766-780.

Fais, L., Kajikawa, S., Amano, S., & Werker, J. F. (2009). Infant Discrimination of a Morphologically Relevant Word‐Final Contrast. *Infancy, 14*(4), 488-499.

Fee, E. J. (1992). Exploring the minimal word in early phonological acquisition. In *Proceedings of the 1992 Annual Conference of the Canadian Linguistics Association*.

Finley, S., & Newport, E. (2010). Morpheme segmentation from distributional information. *Boston University Conference on Language Development (BUCLD) Online Proceedings Supplement.*

Gerken, L., Wilson, R., & Lewis, W. (2005). Infants can use distributional cues to form syntactic categories. *Journal of child language, 32*(2), 249–268. https://doi.org/10.1017/s0305000904006786

Golinkoff, R. M., & Hirsh-Pasek, K. (2006). Introduction: Progress on the verb learning front. In K. Hirsh-Pasek & R. M. Golinkoff (Eds.), *Action meets word: How children learn verbs* (pp. 3–28). Oxford University Press.

Hay, J. (2001). Lexical frequency in morphology: is everything relative? *Linguistics, 39*(6), 1041-1070. https://doi.org/10.1515/ling.2001.041

Hay, J., & Baayen, H. (2002). Parsing and productivity. In Booij, G., Van Marle, J. (Eds.), *Yearbook of Morphology 2001*. Springer.

Hay, J., & Baayen, H. (2003). Phonotactics, parsing and productivity. *Italian Journal of Linguistics, 15*, 99-130.

Hayes, B. (1995) *Metrical stress theory*. Chicago: University of Chicago Press.

Hayes, B. (2011). *Introductory phonology*. John Wiley & Sons.

He, A. X., & Lidz, J. (2017). Verb learning in 14- and 18-month-old English-learning infants. *Language Learning and Development, 13*(3), 335–356. https://doi.org/10.1080/15475441.2017.1285238

Jusczyk, P. W., & Aslin, R. N. (1995). Infants′ detection of the sound patterns of words in fluent speech. *Cognitive psychology, 29*(1), 1-23.

Jusczyk, P. W., Houston, D. M., & Newsome, M. (1999). The beginnings of word segmentation in english-learning infants. *Cognitive psychology, 39*(3-4), 159–207. https://doi.org/10.1006/cogp.1999.0716

Kemler Nelson, D. G., Jusczyk, P. W., Mandel, D. R., Myers, J., Turk, A., & Gerken, L. (1995). The head-turn preference procedure for testing auditory perception. *Infant behavior and development*, 18(1), 111-116.

Kim, Y. J., & Sundara, M. (2021). 6-month-olds are sensitive to English morphology. *Developmental science, 24*(4), e13089. https://doi.org/10.1111/desc.13089

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*(4), 465–494. https://doi.org/10.1006/cogp.1999.0721

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

Mintz T. H. (2003). Frequent frames as a cue for grammatical categories in child directed speech. *Cognition, 90*(1), 91–117. https://doi.org/10.1016/s0010-0277(03)00140-9

Panther, F. A., Mattingley, W., Todd, S., Hay, J. & King, J., (2023) Proto-Lexicon Size and Phonotactic Knowledge are Linked in Non-Māori Speaking New Zealand Adults. *Laboratory Phonology, 14*(1). https://doi.org/10.16995/labphon.7943

Panther, F., Mattingley, W., Hay, J., Todd, S., King, J., & Keegan, P. J. (2024). Morphological segmentations of Non-Māori Speaking New Zealanders match proficient speakers. *Bilingualism: Language and Cognition, 27*(1), 1–15. doi:10.1017/S1366728923000329

Newman, R. S., Sawusch, J. R., & Wunnenberg, T. (2011). Cues and cue interactions in segmenting words in fluent speech. *Journal of Memory and Language, 64*(4), 460-476.

Norris, D., McQueen, J. M., Cutler, A., & Butterfield, S. (1997). The possible word constraint in the segmentation of continuous speech. *Journal of Memory and Language, 44*, 103–132.

Norris, D., McQueen, J. M., Cutler, A., Butterfield, S., & Kearns, R. (2001). Language-universal constraints on speech segmentation. *Language and cognitive processes, 16*(5-6), 637-660.

Oh, Y., Todd, S., Beckner, C., Hay, J., King, J., & Needle, J. (2020). Non-Māori-speaking New Zealanders have a Māori proto-lexicon. *Scientific reports, 10*(1), 22318.

Skoruppa, K., Nevins, A., Gillard, A., & Rosen, S. (2015). The role of vowel phonotactics in native speech segmentation. *Journal of Phonetics, 49*, 67-76.

Sundara, M., Zhou, Z. L., Breiss, C., Katsuda, H., & Steffman, J. (2022). Infants' developing sensitivity to native language phonotactics: A meta-analysis. *Cognition, 221*, 104993.

Vitevitch, M. S., & Luce, P. A. (2004). A web-based interface to calculate phonotactic probability for words and nonwords in English. *Behavior Research Methods, Instruments, & Computers, 36*(3), 481-487.

Willits, J. A., Seidenberg, M. S., & Saffran, J. R. (2014). Distributional structure in language: contributions to noun-verb difficulty differences in infant word recognition. *Cognition, 132*(3), 429–436. https://doi.org/10.1016/j.cognition.2014.05.004